

AD-A123 587

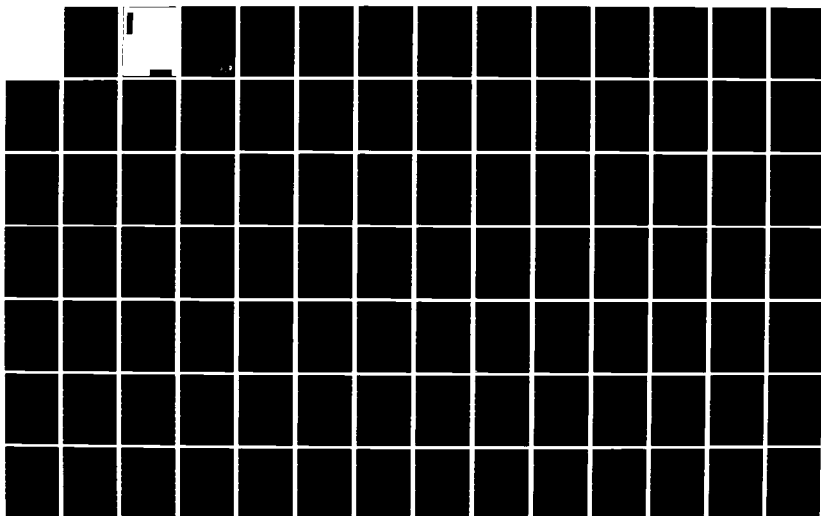
NON-ZERO-SUM TWO-PERSON REPEATED GAMES WITH INCOMPLETE
INFORMATION(U) STANFORD UNIV CA INST FOR MATHEMATICAL
STUDIES IN THE SOCIAL SCIENCES 5 HART FEB 82 TR-367
N00014-79-C-0685

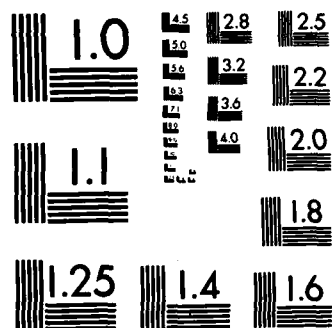
1/2

UNCLASSIFIED

F/G 5/11

NL





MICROCOPY RESOLUTION TEST CHART
NATIONAL BUREAU OF STANDARDS-1963-A

NON-ZERO-SUM TWO-PERSON REPEATED GAMES
WITH INCOMPLETE INFORMATION

by

Sergiu Hart



Technical Report No. 367
February 1982

Accession For	
NTIS GRA&I	<input checked="" type="checkbox"/>
DTIC TAB	<input type="checkbox"/>
Unannounced	<input type="checkbox"/>
Justification	
By _____	
Distribution/	
Availability Codes	
Dist	Avail and/or Special
A	

Prepared Under
NATIONAL SCIENCE FOUNDATION GRANT SES80-06654
AND THE
CENTER FOR RESEARCH ON ORGANIZATIONAL EFFICIENCY
STANFORD UNIVERSITY
Contract ONR-N00014-79-C-0685, United States Office of Naval Research

THE ECONOMICS SERIES
INSTITUTE FOR MATHEMATICAL STUDIES IN THE SOCIAL SCIENCES
Fourth Floor, Encina Hall
Stanford University
Stanford, California
94305

DTIC
ELECTE
S JAN 21 1983 D

This document has been approved
for public release and sale; its
distribution is unlimited.

p-1

NON-ZERO-SUM TWO-PERSON REPEATED GAMES
WITH INCOMPLETE INFORMATION*

by

Sergiu Hart**

1. Introduction

An incomplete information environment is one where at least some of the participants do not possess all the relevant data. Much interest has been devoted in recent years to the analysis of such situations. In the economic theory literature, for example: the principal-agent problem; the theory of auctions; signalling (e.g., in insurance markets); rational expectations equilibria; and so on.

What are the main difficulties in such problems? First, consider the "informed" persons--those who know more than others. On one hand, it is to their advantage to make use of their additional information (in order to improve their own final outcome). On the other hand, by doing so they actually reveal this information--and their relative advantage vanishes. Thus--what is the good of being more informed, if one cannot profit from it? This type of conflict is an essential issue in the analysis of incomplete information environments.

*This work was partially supported by National Science Foundation Grant SES80-06654 and Contract ONR-N00014-79-C-0685 at the Institute for Mathematical Studies in the Social Sciences, Stanford University. It was also supported by CORE and the Mathematics Department, Université Catholique de Louvain.

I wish to acknowledge many useful discussions with R.J. Aumann, J.-F. Mertens and S. Zamir.

**Department of Statistics, Tel-Aviv University, Tel-Aviv, Israel 69978.

As an idealized example, assume someone has "inside information" that a certain small company has just succeeded in developing a new product, for which a very profitable market exists. He thus expects that the value of the shares of this company in the Stock Exchange will raise dramatically. Should he immediately buy a large quantity of these shares? By doing so, he will implicitly signal to the others the success of the company--and everyone will want to buy its shares, raising their value immediately and lowering the profits of the initially informed person. The answer clearly lies in him buying the "right" quantity of shares--not too large to draw attention, and not too small to make his profit insignificant.

The results of the analysis of such models of incomplete information usually indicate that some transmission of information does occur (possibly, in an implicit way only; namely, deducing information from actions taken by those possessing it). Thus, there is need for communications, and some sort of cooperation may arise (e.g., "trading information")--even though everything is based on purely selfish (non-cooperative) motives.

There is yet another conflict--this time, for the "uninformed" participants. Should they trust the information transmitted by the informed ones? In the Stock Exchange example--maybe the purpose of buying a large quantity of shares is just to convince everyone that a technological breakthrough indeed occurred, leading to a big buying activity, which may finally make a good profit for the one that started

it all. This, also in case no new product has been at all developed by the company!

Game theory is a tool for studying conflict situations--by definition, inter-personal conflicts. However, one obtains as an outcome resolution of intra-personal conflicts (like the ones mentioned above) as well--based on individual rational behaviour. This is true in particular for games with incomplete information--a class of which forms the subject of this paper.

An important development in game theory in recent years has been in the study of multi-stage games--especially, the so called repeated games, where the same game is played repeatedly. This suggests itself as a good framework for incomplete information games, for two main reasons.

The first one is that by its very nature, a repeated game has enough structure to allow the kinds of complicated behaviour we described above (and many others as well). There is enough "time" to enable players to "generate" certain beliefs in other people, or to make deductions, statistical inferences, and so on. There is also place for threats, for punishments--and for rewards too.

The second reason is more formal--although closely related to the first one. Consider an infinitely repeated game with complete information. A well known result (called the "Folk Theorem" since its authorship is not clear) states that the non-cooperative equilibria in

the repeated game precisely correspond to the individually rational and jointly feasible points in the one-shot game. The importance of this result is that one obtains cooperative outcomes in the one-shot game from non-cooperative behaviour in the infinite game. Thus, the cooperation we usually observe is explained here not as an outcome of altruistic motives--but of purely selfish non-cooperative ones (which many feel are the only rational ones).

One is therefore led in a natural way to the study of repeated games of incomplete information. The first research on these was done in the *Mathematica* [1966-68] reports, in particular by Aumann, Maschler and Stearns. It turned out that the very complex structure of these games--which, as we pointed out above, is one of the reasons for studying them--creates many difficulties. Up to date, essentially only two-person zero-sum games have been completely analyzed (see the forthcoming book of Mertens and Zamir [1980], or the notes of Sorin [1980] for details).

As for the non-zero-sum case (still, only two players), a first study has been done by Aumann, Maschler and Stearns [1968]. They characterized a special class of equilibria, in the so-called standard one-sided information case, where one player has more information than the other one, and both observe during the play all the actions taken. These equilibria--called "enforceable joint plans"--essentially consist of a transmission of information from the informed to the uninformed player ("signalling"), followed by a completely non-revealing play from then on (similar to the Folk Theorem). Moreover,

they showed that this does not exhaust all equilibria--one could have joint randomizations of enforceable joint plans, and so on.

Our main result in this paper is the complete characterization of all equilibria in such games. We will show that every equilibrium is equivalent to a collection of non-revealing "plans", one of which is chosen at random. This choice is done via a sequence of communications, which are of two types: signalling (i.e., implicit transmission of information), and jointly controlled randomizations (i.e., "lotteries" in which no one player can unilaterally change the probabilities).^{1/}

Thus, we are able to characterize in a formal way all the kinds of cooperation and communication that arise out of non-cooperative behaviour in these games; moreover, we obtain a precise structure that guarantees it does not pay any player to do anything else (e.g., revealing less or more, or double-crossing, cheating, and so on). We would like to point out that the model is not the most general possible (in particular, in terms of the information structure); this paper is to be regarded as a first step in the analysis of non-zero-sum repeated games with incomplete information.

The formal model is described in Section 2, together with various notions of equilibrium. The main results are stated in Section 3, which also includes additional discussion and intuitive interpretations. Sections 4 and 5 are devoted to the two parts of the proof, and in Section 6 we present some results on enforceable joint plans. We would like to point out that Sorin [1981] has recently proved the existence of such equilibria whenever the number of possible games is two.

Some notation : R is the real line, and R^n the n -dimensional Euclidean space. For vectors $x = (x_1, \dots, x_n)$ and $y = (y_1, \dots, y_n)$ in R^n , $x \geq y$ means $x_i \geq y_i$ for all $i = 1, 2, \dots, n$, and $x \cdot y$ is the scalar product $\sum_{i=1}^n x_i y_i$. For a finite set L , $|L|$ is the number of elements of L , and R^L the $|L|$ -dimensional Euclidean space with coordinates indexed by the members of L (thus, we write $x = (x_\ell)_{\ell \in L} = (x(\ell))_{\ell \in L}$ for x in R^L). The unit simplex in R^L will be denoted by Δ^L :

$$\Delta^L = \{x \in R^L : x_\ell \geq 0 \text{ for all } \ell \text{ in } L, \sum_{\ell \in L} x_\ell = 1\}.$$

Finally, N is the set of positive integers $\{1, 2, \dots\}$.

2. The Model

The class of games we study is given by the following:

- (i) Two players, player 1 and player 2.
- (ii) A finite set I of choices for player 1 and a finite set J of choices for player 2; I and J contain each at least two elements 2/
- (iii) A finite set K of games; to each k in K there corresponds a pair of $I \times J$ matrices (A^k, B^k) , with

$$A^k = (A^k(i, j))_{\substack{i \in I \\ j \in J}}, \quad B^k = (B^k(i, j))_{\substack{i \in I \\ j \in J}}.$$

(iv) A probability vector $p = (p^k)_{k \in K}$ on the set K (i.e., $p \in \Delta^K$; without loss of generality, we assume $p^k > 0$ for all k in K - otherwise, we may discard those k that have zero probability).

Based on (i)-(iv), a game of incomplete information $\Gamma_\infty(p)$ is given as follows:

(v) An element κ of K is chosen according to the probability vector p ; player 1 is told κ , player 2 is not.

(vi) At each stage $t = 1, 2, \dots$, player 1 chooses an element i_t in I and player 2 chooses an element j_t in J ; the choices are made simultaneously (or, without either player knowing what the other did).

(vii) Both players are then told the pair (i_t, j_t) , and they get the payoffs $A^k(i_t, j_t)$ and $B^k(i_t, j_t)$, respectively (but they do not observe these payoffs).

(viii) Both players have perfect recall (i.e., they do not forget what they were told at all previous stages).

(ix) All of (i)-(viii) is common knowledge to both players (see Aumann [1976] for a precise definition).

Usually, (v) is called one-sided information (see also the discussion below), and (vii) and (viii) - standard information. Note that the players observe only the actual choices i_t and j_t , and not the randomizations used.

Following Harsanyi [1967-68], this can be equivalently viewed as a game with complete but imperfect information (namely, where the uncertainty players have is not about the "rules of the game" - e.g.,

payoffs - but only on moves previously made, by the players or by chance). This is done by adding a stage $t = 0$, at which "nature" chooses an element κ of K according to the probability p . At each stage $t = 1, 2, \dots$, the information player 2 has consists of the sequence of previous choices by both players: $(i_1, j_1), (i_2, j_2), \dots, (i_{t-1}, j_{t-1})$. As for player 1, he in addition knows κ .

This completes the description of $\Gamma_\infty(p)$. It should be pointed out that more general games can be made to fit into this model. In particular, consider the case where player 1 does not have full information on κ , but player 2 knows even less (see Mertens & Zamir [1980, Ch.III]). Formally, a partition of the set K is given for each player, which is informed only what element of his partition contains the chosen κ . For example, let $K = \{1, 2, 3, 4, 5\}$, the partition of player 1 is $\{1, 2\}, \{3\}, \{4\}, \{5\}$, and that of player 2 is $\{1, 2, 3\}, \{4, 5\}$. The (common) prior is $p = (1/5, 1/5, 1/5, 1/5, 1/5)$. First, we observe that both players distinguish between $\{1, 2, 3\}$ and $\{4, 5\}$ - thus, there are two completely disjoint games. In the first, both do not distinguish between 1 and 2; therefore, this corresponds to $K' = \{\{1, 2\}, \{3\}\}$ and $p' = (2/3, 1/3)$, where the payoff matrix for $\{1, 2\}$ is $A^{\{1, 2\}} = (1/2)A^1 + (1/2)A^2$ and similarly for B . Note that $2/3$ is the conditional probability that $\kappa \in \{1, 2\}$ given $\kappa \in \{1, 2, 3\}$, $1/3$ is $P(\kappa = 3 | \kappa \in \{1, 2, 3\})$, and $1/2 = P(\kappa = 1 | \kappa \in \{1, 2\}) = P(\kappa = 2 | \kappa \in \{1, 2\})$. In the second latter game, $K'' = \{4, 5\}$ and $p'' = (1/2, 1/2)$. Thus, the original game has been decomposed into two games, each fitting our model. It should be clear

how to generalize the construction given for this example.

Next we describe the sets of strategies of the players in $\Gamma_\infty(p)$. For each $t = 1, 2, \dots$, let H_t be the set of histories up to (but not including) stage t , namely,^{3/}

$$H_t = (I \times J)^{t-1}.$$

A pure strategy σ of player 1 is a collection $\sigma = \{\sigma_t\}_{t=1}^\infty$, where

$$(2.1) \quad \sigma_t: H_t \times K \rightarrow I$$

for all $t = 1, 2, \dots$. Thus, for every history h_t in H_t and every k in K (the "true" game κ chosen), $\sigma_t(h_t; k)$ is the choice i_t made by player 1. In a similar way, a pure strategy τ of player 2 is $\tau = \{\tau_t\}_{t=1}^\infty$, where

$$(2.2) \quad \tau_t: H_t \rightarrow J$$

for all $t = 1, 2, \dots$.

A mixed strategy is, as usual, a probability distribution over the set of pure strategies. Since $\Gamma_\infty(p)$ is a game with perfect recall, one can restrict the study to behaviour strategies (cf. Kuhn [1953] and Aumann [1964]), where players make independent randomizations at each move. A behaviour strategy is thus defined in the same way as a pure strategy, with (2.1) replaced by

$$(2.3) \quad \sigma_t: H_t \times K \rightarrow \Delta^I,$$

and (2.2) replaced by

$$(2.4) \quad \tau_t: H_t \rightarrow \Delta^J.$$

Since we never use pure strategies specifically, the term "strategy" will henceforth mean behaviour or mixed strategy.

We have not yet defined payoffs in $\Gamma_\infty(p)$ - only sequences of payoffs. Given a pair of strategies (σ, τ) of the two players, we denote

$$(2.5) \quad a_T^k = \frac{1}{T} \sum_{t=1}^T A^k(i_t, j_t),$$

$$(2.6) \quad \beta_T = \frac{1}{T} \sum_{t=1}^T B^k(i_t, j_t),$$

for all $T = 1, 2, \dots$ and all k in K . Thus, a_T^k is the average payoff up to (and including) stage T to player 1, if the true game is $\kappa = k$; this depends on the choices of i_t 's and j_t 's, made according to σ and τ (actually, only $\sigma(\cdot; k)$ and τ matter). Let $E_{\sigma, \tau}^k(a_T^k)$ denote its expectation. For player 2, β_T is his average payoff up to T ; it depends on σ, τ and also on the choice of κ (according to p). Let $E_{\sigma, \tau, p}(\beta_T)$ be its expectation.

A pair (σ, τ) of strategies is a (Nash) equilibrium point in $\Gamma_\infty(p)$ if

$$(2.7) \quad \liminf_{T \rightarrow \infty} E_{\sigma, \tau}^k(a_T^k) \geq \limsup_{T \rightarrow \infty} E_{\sigma', \tau}^k(a_T^k)$$

for all strategies σ' of player 1 and all k in K , and

$$(2.8) \quad \liminf_{T \rightarrow \infty} E_{\sigma, \tau, p}(\beta_T) \geq \limsup_{T \rightarrow \infty} E_{\sigma, \tau', p}(\beta_T)$$

for all strategies τ' of player 2. If we take $\sigma' = \sigma$ in (2.7), we get a vector $a = (a^k)_{k \in K}$ such that

$$(2.9) \quad \lim_{T \rightarrow \infty} E_{\sigma, \tau}^k(a_T^k) = a^k$$

for all k in K . Similarly, $\tau' = \tau$ in (2.8) gives β with

$$(2.10) \quad \lim_{T \rightarrow \infty} E_{\sigma, \tau, p}(\beta_T) = \beta.$$

We will call a and β the payoffs of the equilibrium point (σ, τ) .

Note that they are computed ex-post--namely, after the choice of κ was made and player 1 was informed of it. Therefore, player 1 considers his payoffs in each possible state $\kappa = k$, whereas for player 2 only his expectation over κ matters. It can be easily checked that the definition does not change if we replace (2.7) by ex-ante optimality, namely:

$$\liminf_{T \rightarrow \infty} E_{\sigma, \tau, p}(\alpha_T) \geq \limsup_{T \rightarrow \infty} E_{\sigma', \tau, p}(\alpha_T),$$

where α_T is defined in the same way as β_T (thus, $\alpha_T = a_T^k$). Indeed, since the value k of κ is in any case part of the information player 1 has at every stage, he can choose his best response against

τ independently for each k . In the imperfect information version of $\Gamma_\infty(p)$, the adequate payoff is indeed this expectation over κ ; in the incomplete information one, the vector of payoffs should be considered instead since, given any "type" k - in Harsanyi's terminology - it does not care about the payoffs to all the other possible types!

A strengthening of the definition of equilibrium is suggested by the results obtained in the zero - sum case (i.e., where $A^k + B^k \equiv 0$ for all k in K). A pair of strategies (σ, τ) is a uniform equilibrium point in $\Gamma_\infty(p)$ if

$$(2.11) \quad \liminf_{T \rightarrow \infty} E_{\sigma, \tau}^k(a_T^k) \geq \limsup_{T \rightarrow \infty} (\sup_{\sigma'} E_{\sigma', \tau}^k(a_T^k))$$

for all k in K , and

$$(2.12) \quad \liminf_{T \rightarrow \infty} E_{\sigma, \tau, p}(\beta_T) \geq \limsup_{T \rightarrow \infty} (\sup_{\tau'} E_{\sigma, \tau', p}(\beta_T)) .$$

Clearly, every uniform equilibrium point is also an equilibrium point (if we change the order of limsup and sup in (2.11) and (2.12), we obtain (2.7) and (2.8), respectively). The payoffs (a, β) are given by (2.9) and (2.10).

To emphasize the difference between the two definitions, we translate them into the " ϵ - language". A uniform equilibrium satisfies the following: for every $\epsilon > 0$ there exists $T_0 \equiv T_0(\epsilon)$ large enough such that for all $T > T_0$,

$$(2.13) \quad E_{\sigma', \tau}^k(a_T^k) \leq a^k + \epsilon \quad \text{and}$$

$$E_{\sigma, \tau', p}(\beta_T) \leq \beta + \varepsilon ,$$

for all k in K and all strategies σ' of player 1 and τ' of player 2. For a regular equilibrium (according to (2.7) and (2.8)), T_0 may also depend on σ' and τ' . The importance of (2.13) uniformly in σ' and τ' is that it implies that (σ, τ) generates an ε - equilibrium in all long enough but finite games $\Gamma_T(p)$ (which are defined in the same way as $\Gamma_\infty(p)$, but they only last T stages). Since $\Gamma_\infty(p)$ may be viewed as an "idealization" of such games, the uniform definition may seem more appropriate.

However, we will prove the following result:

Proposition 2.14: The sets of payoffs of equilibrium points and of uniform equilibrium points in $\Gamma_\infty(p)$, coincide.

Thus, although it is clear that there exist equilibrium points that are not uniform, they are always payoff-equivalent to uniform ones.

Other definitions of equilibrium are also possible. For example, one could use Abel instead of Cesaro summability; namely, limits as $\rho > 0$ converges to 0 of

$$E(\rho \sum_{t=1}^{\infty} \frac{x_t}{(1+\rho)^t}) ,$$

where $\{x_t\}_{t=1}^{\infty}$ is the corresponding sequence of payoffs (this is interpreted as the limit, as the interest rate goes to zero, of the current value). Banach limits (see Section 4) can also be used.

However, in all cases the set of equilibrium payoffs will not change.

In view of this result, we can unambiguously define the set of equilibrium payoffs of $\Gamma_{\infty}(p)$. Our main result will be a characterization of this set.

Can one further strengthen the definition of equilibrium by changing the order of limit and expectation? The answer is no - as an example by J.-F. Mertens and the author shows already in the zero-sum case.

3. Statement and Interpretation of the Main Result

In this section we state our main result - the characterization of all equilibria in $\Gamma_{\infty}(p)$.

The Folk Theorem in the complete information case states that the set of equilibrium payoffs coincides with the set of feasible and individually rational payoffs. We consider first the notion of individual rationality; it is to be understood in the sense of what each player cannot be prevented from obtaining (i.e., the "minmax"). The study of the zero-sum case (Aumann & Maschler [1966]) enables us to characterize individual rationality in $\Gamma_{\infty}(p)$.

We need some notation first. Let p be a probability vector in Δ^K ; let $p \cdot A$ be the matrix $\sum_{k \in K} p^k A^k$ (i.e., whose (i,j) th element is $\sum_{k \in K} p^k A^k(i,j)$). Consider the two-person zero-sum game with payoffs to player 1 given by $p \cdot A$, and let $(val_1 A)(p)$ denote its value (when played just once). Thus,

$$(3.1) \quad (val_1 A)(p) = \max_{x \in \Delta^I} \min_{y \in \Delta^J} (p \cdot A)(x, y) = \min_{y \in \Delta^J} \max_{x \in \Delta^I} (p \cdot A)(x, y) ,$$

where $x = (x_i)_{i \in I}$, $y = (y_j)_{j \in J}$, and

$$(p \cdot A)(x, y) = \sum_{i \in I} \sum_{j \in J} x_i y_j \sum_{k \in K} p^k A^k(i, j) .$$

Similarly, let $(\text{val}_2 B)(p)$ be the value to player 2 of the two-person zero-sum game with payoff matrix $p \cdot B$ to player 2. Clearly,

$$(3.2) \quad (\text{val}_2 B)(p) = - (\text{val}_1 (-B))(p) .$$

For a function f on Δ^K , let $\text{vex } f$ denote its convexification; namely, $\text{vex } f$ is the largest convex function on Δ^K that does not exceed f . We will write $(\text{vex } \text{val}_2 B)(p)$ for the evaluation of the function $\text{vex } (\text{val}_2 B)$ at the point p .

We can now define: a vector $a = (a^k)_{k \in K}$ in R^K is an individually rational payoff vector to player 1 in $\Gamma_\infty(p)$ if

$$(3.3) \quad q \cdot a \geq (\text{val}_1 A)(q) , \text{ for all } q \text{ in } \Delta^K .$$

A scalar β in R is an individually rational payoff to player 2 in $\Gamma_\infty(p)$ if

$$(3.4) \quad \beta \geq (\text{vex } \text{val}_2 B)(p) .$$

These definitions are the correct ones, in view of the following results. A set Q in R^K is approachable by player 2 (cf. Blackwell [1956]) if there exists a strategy τ of player 2 such that

$$\lim_{T \rightarrow \infty} \left(\sup_{\sigma} E_{\sigma, \tau} (d(Q, a_T)) \right) = 0 ,$$

where $a_T = (a_T^k)_{k \in K}$ (recall (2.5)), d is the Euclidean distance in R^K , and the supremum is over all strategies σ of player 1.

Proposition 3.5: Let a be a vector in R^K . Then (3.3) is a necessary and sufficient condition for the set $Q = \{x \in R^K : x \leq a\}$ to be approachable by player 2.

Proof: Blackwell [1956]; for example, see Aumann & Maschler [1966]. Q.E.D.

Thus, if (3.3) is satisfied, then player 2 can guarantee that the payoffs to player 1 will not, in the limit, exceed a^k for all k in K simultaneously. If (3.3) is not satisfied, then given any strategy of player 2, player 1 has a strategy such that, for at least one k in K , he will get more than a^k .

For player 2, we have

Proposition 3.6: Let β be a scalar in R . Then (3.4) is a necessary and sufficient condition for player 1 to have a strategy σ such that

$$\limsup_{T \rightarrow \infty} \left(\sup_{\tau} E_{\sigma, \tau, p}(\beta_T) \right) \leq \beta ,$$

where β_T is given by (2.6) and the supremum is over all strategies τ of player 2.

Proof: Aumann & Maschler [1966]; see (3.2).

Q.E.D.

Again, this means that player 1 can hold player 2 down to $(\text{vex val}_2 B)(p)$ in $\Gamma_\infty(p)$, but to no less than that.

Having completed the study of individual rationality, we come next to feasibility. Let us consider a simple case first. Fix i in I and j in J : is there an equilibrium resulting in the pair (i,j) being chosen at every stage? Clearly, the answer depends on the actions the players will take "outside of equilibrium" - namely, when (i,j) is no longer been played. Again as in the Folk Theorem, it is easy to see that the necessary and sufficient condition is precisely individual rationality for both players (each will use the corresponding strategy given by Propositions 3.5 and 3.6, respectively, immediately after the other deviates from (i,j)). Therefore, the payoffs $\alpha = (A^k(i,j))_{k \in K}$ and $\beta = \int_{k \in K} p^k B^k(i,j)$ will be equilibrium payoffs in $\Gamma_\infty(p)$ if and only if (3.3) and (3.4) are satisfied.

This reasoning can now be extended to any convex combination by using the corresponding frequencies. It generates a class of equilibria in $\Gamma_\infty(p)$, which result in player 1 actually playing the same for all k in K (i.e., independent of κ). Note that this is true only "in equilibrium" (i.e., so long as there are no defections); "out of equilibrium", the strategy given by Proposition 3.6 may depend on κ . We will thus call these equilibria non-revealing.

To define formally the corresponding payoffs, we denote by $(A,B)(i,j)$ the vector

$$(A,B)(i,j) = ((A^k(i,j))_{k \in K}, (B^k(i,j))_{k \in K}) \in R^K \times R^K ,$$

for all i in I and j in J . Then, let

$$(3.7) \quad F = \text{conv} \{(A,B)(i,j) : i \in I, j \in J\} ,$$

where "conv" denotes the convex hull of a set. F can be viewed as the set of feasible vector payoffs (in the one-shot game).

Let M be the maximum absolute value of any possible payoff:

$$(3.8) \quad M = \max \{|A^k(i,j)|, |B^k(i,j)| : i \in I, j \in J, k \in K\} .$$

We then write R_M^K for the set of all vectors in R^K , all of whose coordinates are bounded by M . We also put R_M for the real interval $[-M,M]$ (thus $R_M^K = (R_M)^K$). Clearly, F is a subset of $R_M^K \times R_M^K$.

Finally, we define the set G as follows: it consists of all triples (a,β,p) , with a in R_M^K , β in R_M and p in Δ^K , such that (3.3) and (3.4) are satisfied, and there exist c and d in R^K with

$$(3.9) \quad (c,d) \in F ,$$

$$(3.10) \quad a \geq c \text{ and } p \cdot a = p \cdot c ,$$

$$(3.11) \quad \beta = p \cdot d .$$

As in the zero-sum case, we will find it necessary to consider all the games $\Gamma_\infty(p)$, as p ranges over Δ^K , at the same time; a triple (a,β,p) is understood as (a,β) being payoffs in $\Gamma_\infty(p)$.

In view of our previous discussion, G is essentially the set of

payoffs corresponding to non-revealing equilibria (note that (3.10) can be restated as: $a^k = c^k$ if $p^k > 0$, $a^k \geq c^k$ otherwise - therefore, a and c are identical for all relevant games).

Our main result states that, based on the set G , we can characterize all equilibrium payoffs. We thus define the concept of a G -process, as follows.

Let $g = (a, \beta, p) \in R_M^K \times R_M \times \Delta^K$. A sequence $\{g_n\}_{n=1}^\infty = \{(a_n, \beta_n, p_n)\}_{n=1}^\infty$ of $(R_M^K \times R_M \times \Delta^K)$ - valued random variables (on some probability space) is called a G -process starting at g if:

$$(3.12) \quad g_1 \equiv g \text{ a.s. .}$$

$$(3.13) \quad \text{There exists a non-decreasing sequence } \{Z_n\}_{n=1}^\infty \text{ of finite fields}^{4/} \text{ with respect to which } \{g_n\}_{n=1}^\infty \text{ is a martingale.}$$

$$(3.14) \quad \text{Let } g_\infty \text{ be an a.s. limit of } g_n \text{ (as } n \rightarrow \infty); \text{ then}$$

$$g_\infty \in G \text{ a.s. .}$$

$$(3.15) \quad \text{For each } n = 1, 2, \dots, \text{ either } a_{n+1} = a_n \text{ a.s., or}$$

$$p_{n+1} = p_n \text{ a.s.}$$

The martingale condition in (3.13) means that g_n is Z_n -measurable and $E(g_{n+1} | Z_n) = g_n$ a.s. for all n . Together with (3.12), it implies

$E(g_n) = g$ for all n . Since the sequence is uniformly bounded, the Martingale Convergence Theorem implies that it has an a.s. limit - thus (3.14) is well defined. It then means that $g_\infty = (a_\infty, \beta_\infty, p_\infty)$ satisfies

a.s. individual rationality for both players (i.e., (3.3) and (3.4)), and also (3.9) - (3.11).

The last condition (3.15) is slightly unusual; at every step, either a_n or p_n remain constant (while the other may change - but in such a way that the conditional expectation does not, by (3.13)). If we disregard the β_n coordinate, such a process may be called a bi-martingale (see Proposition 3.18 below). The study of such objects will be taken up in a forthcoming paper of R. J. Aumann and the author.

Finally, we define G^* as the set of all points $g = (a, \beta, p)$ in $R_M^K \times R_M \times \Delta^K$ such that there exists a G-process starting at g . We note here that (3.12) and (3.15) are essential conditions; without either one, G^* will just be the convex hull of G .

We are now ready to state our main result.

Main Theorem: Let $a \in R^K$ and $\beta \in R$. Then (a, β) are equilibrium payoffs in $\Gamma_\infty(p)$ if and only if $(a, \beta, p) \in G^*$.

Thus, the set G^* is the graph of the equilibrium payoffs correspondence (as p ranges over Δ^K).

The Main Theorem and Proposition 2.14 will be proved together (we know of no direct proof of the latter alone). This will be done by showing first (in Section 4) that all equilibrium payoffs, according to the regular definition (2.7) - (2.10), belong to G^* . And second, by constructing (in Section 5) a uniform equilibrium (cf. (2.11) and (2.12)) corresponding to any point in G^* .

The second part of the proof leads us to an important additional result; namely, that all equilibrium points in $\Gamma_\infty(p)$ are equivalent to a special class of equilibria (those we construct in Section 5). Informally, such an equilibrium consists of a "master plan", which is followed by each player so long as the other does it too; and of "punishments", which come into effect after a deviation from the master plan has been detected.

The master plan is a sequence of "communications" between the two players, the purpose of which is to eventually settle on a point in G which is played forever from then on (using frequencies), and leads to the desired "payoffs". The communications are of two sorts: "signalling", where the informed player 1 plays dependent on κ (and thus reveals some of his information to player 2, who can update his posterior probabilities); and joint decisions, more precisely "jointly controlled lotteries", where the two players make together a randomization on how to continue the play. Signalling has already been obtained in the zero-sum case; however, the jointly controlled lotteries (in which the uninformed player plays a no lesser role than the informed one) are a feature of the non-zero-sum case only.

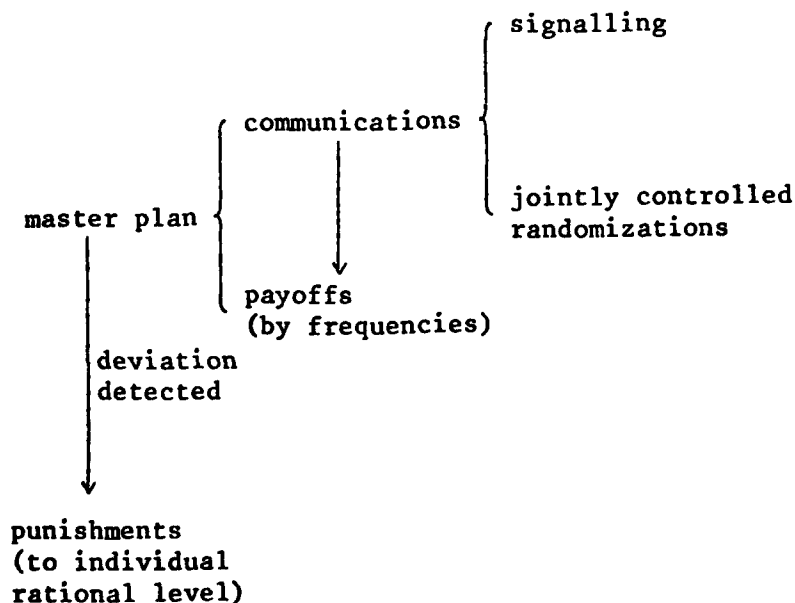
At the end of the communication period (which we assume for the moment to consist of finitely many stages only), player 1 will play independent of κ (otherwise, he will reveal additional information) - and thus a non-revealing equilibrium results from then on (a point in G). In the general case, the sequence of communications may be infinite. However, after a long enough time, almost everything that was

ever going to be revealed (by player 1) or decided (by a joint randomization) has already occurred - and we are essentially at a non-revealing point again (i.e., in G). To generate the right payoffs, "payoff accumulation" periods are then introduced between communications - at which both players choose prescribed moves (again - with the correct frequencies).

Finally, punishments are always in accordance to the strategies given by Propositions 3.5 and 3.6, respectively (see Proposition 3.16).

The structure of such equilibria is summarized in Figure 1.

Figure 1



The G-process is thus "followed" during the play. At each stage, the corresponding $g_n = (a_n, \beta_n, p_n)$ will serve as a "state variable", with a_n in R_M^K being the vector payoff player 1 will get from then on, β_n in R_M the same for player 2 (averaging over k), and p_n in Δ^K the vector of posterior probabilities for κ .

Why is an equilibrium thus obtained? Deviations during communications stages are not helpful: jointly controlled lotteries are so designed as to have each one of the players generate the right probabilities even if the other does not; as for signalling by player 1, it occurs precisely when $a_{n+1} = a_n$ in the G-process, which makes him indifferent among the various alternatives. In all other cases, the punishments keep the players in line. This is due to the following:

Proposition 3.16: Let $\{(a_n, \beta_n, p_n)\}_{n=1}^{\infty}$ be an $(R_M^K \times R_M \times \Delta^K)$ -valued martingale, converging a.s. to $(a_{\infty}, \beta_{\infty}, p_{\infty})$. Then:

- (i) a_{∞} satisfies (3.3) a.s. if and only if a_n satisfy (3.3) a.s. for all $n = 1, 2, \dots$.
- (ii) $(\beta_{\infty}, p_{\infty})$ satisfies ^{6/} (3.4) a.s. if and only if (β_n, p_n) satisfy (3.4) a.s. for all $n = 1, 2, \dots$.

Proof: The "if" part is obtained by taking the limit as $n \rightarrow \infty$ (in (ii), we use the continuity of the function vex val_2^B - e.g., see Mertens & Zamir [1980, Theorem 3.14]).

Let $\{Z_n\}_{n=1}^{\infty}$ be the corresponding sequence of σ -fields, then we have $a_n = E(a_{\infty} | Z_n)$ by the martingale theorem. The "only if" part

in (i) is obtained by taking conditional expectations over Z_n . As for (ii),

$$\begin{aligned}\beta_n &= E(\beta_\infty | Z_n) \geq E((\text{vex val}_2 B)(p_\infty) | Z_n) \\ &\geq (\text{vex val}_2 B)(E(p_\infty | Z_n)) = (\text{vex val}_2 B)(p_n) ,\end{aligned}$$

where we used the convexity of the function $\text{vex val}_2 B$. Q.E.D.

This last result leads to an additional interpretation of G^* as outcomes of bargaining processes - see Aumann [1981].

Corollary 3.17: Let (a, β) be equilibrium payoffs in $\Gamma_\infty(p)$. Then a and β are individually rational for player 1 and player 2, respectively.

Proof: Proposition 3.16 for $n = 1$ (recall (3.12)). Q.E.D.

Another property of a G -process (which led to the name "bi-martingale") is as follows.

Proposition 3.18: Let $\{(a_n, p_n)\}_{n=1}^\infty$ be an $(R_M^K \times \Delta^K)$ -valued martingale with respect to a non-decreasing sequence of σ -fields $\{Z_n\}_{n=1}^\infty$. If (3.15) is satisfied, then $\{a_n \cdot p_n\}_{n=1}^\infty$ is also a martingale with respect to $\{Z_n\}_{n=1}^\infty$.

Proof: Let n be such that $a_{n+1} = a_n$ a.s.; then $E(a_{n+1} \cdot p_{n+1} | Z_n) = E(a_n \cdot p_{n+1} | Z_n) = a_n \cdot E(p_{n+1} | Z_n) = a_n \cdot p_n$, since a_n is Z_n -measurable. The same when $p_{n+1} = p_n$. Q.E.D.

4. From Equilibrium to Martingale

This section contains the proof of the first part of our result; namely, given an equilibrium point we construct the corresponding G-process (see Proposition 4.43 at the end of the section for a precise statement).

We start with an informal discussion of the proof. Let (σ, τ) be an equilibrium point; to simplify the arguments, let us assume that the frequencies with which the various pairs (i, j) in $I \times J$ are played, always converge. Let $c_\infty = (c_\infty^k)_{k \in K}$ and $d_\infty = (d_\infty^k)_{k \in K}$ be the limit payoffs, then clearly $(c_\infty, d_\infty) \in F$. For every history h_t up to stage t , we then define the following: for each k in K , $a^k(h_t)$ is the expected payoff to player 1 if $\kappa = k$ (thus, $a^k(h_t)$ is just the expectation of c_∞^k given h_t); $\delta(h_t)$ is the expected payoff to player 2 (the expectation of d_∞^k given h_t); and $p^k(h_t)$ is the (posterior) probability that $\kappa = k$ (again, given h_t). We next introduce "half-steps", i.e., we define the above conditional expectations when given both h_t and the next move i_t of player 1; we will thus write $a^k(h_t, i_t)$, and so on.

Assume h_t has positive probability of occurring when $\kappa = k$. Then all possible moves i_t of player 1 (i.e., those with $\sigma(h_t; k)(i_t) > 0$) must have the same expected payoff $a^k(h_t, i_t)$. Otherwise, player 1 could give probability 1 to that i_t leading to the highest payoff; this would be "undetected" by player 2 (since this i_t is possible according to σ), thus giving an expected payoff

of $a^k(h_t, i_t)$, which is higher than that given by σ . This contradicts the equilibrium conditions, therefore $a^k(h_t, i_t)$ must be constant--hence, equal to $a^k(h_t)$ --for all possible i_t 's. A similar argument shows that for the other i_t 's an inequality is obtained (if they are not chosen, then their corresponding payoff cannot be higher); this will eventually lead to the condition (3.10).

Next, consider the half-step from (h_t, i_t) to $h_{t+1} = (h_t, i_t, j_t)$. Since player 2 does not know κ , j_t is independent of it, and the posterior probabilities cannot change. We thus have $p^k(h_t, i_t) = p^k(h_{t+1})$.

It is easy to check that in all other cases, the martingale conditions are satisfied; e.g., $E(a^k(h_{t+1}) | h_t, i_t) = a^k(h_t, i_t)$, and so on. We have therefore obtained a martingale (with the index set being that of half-steps), which furthermore satisfies (3.15). The individual rationality conditions (3.3) and (3.4) also hold (since otherwise (σ, τ) will not be an equilibrium), and one can show that, in the limit (which exists by the martingale convergence theorem), a point in G is a.s. reached.

The actual proof will be quite complicated. Since we have no convergence of the payoffs, we will need to use Banach limits. To facilitate following the arguments, we divided the proof into a sequence of subsections.

4.1 The Probability Space

For each $t \in \mathbb{N}$ (the set of positive integers), we defined

$$H_t = (I \times J)^{t-1},$$

the set of histories before stage t . We also define the set of infinite histories

$$H_\infty = \prod_{t=1}^{\infty} (I \times J),$$

an element of H_∞ being a sequence $\{(i_t, j_t)\}_{t=1}^{\infty}$ of moves made by the two players at all stages.

On H_∞ we define for each $t \in \mathbb{N}$ the finite field generated by H_t , and call it \mathcal{H}_t ; thus, two infinite histories belong to the same atom in \mathcal{H}_t if and only if they coincide up to (but not including) t . Let \mathcal{H}_∞ be the σ -field generated by all the \mathcal{H}_t 's (usually called the cylindrical or the product σ -field on the space H_∞).

The basic probability space will also include the choice of κ in K by chance. Thus, let $\Omega = H_\infty \times K$ be endowed with the σ -field $\mathcal{H}_\infty \otimes 2^K$. Each pair of strategies (σ, τ) and each probability vector $p \in \Delta^K$ for the initial chance move determine a probability distribution on this space. We denote it by $P_{\sigma, \tau, p}$; note that $E_{\sigma, \tau, p}$ used in Section 2 is precisely the expectation with respect to $P_{\sigma, \tau, p}$, and $E_{\sigma, \tau}^k$ is the conditional expectation given $\kappa = k$.

We will use some additional fields on H_∞ . For each $t \in \mathbb{N}$, let

$$H_{t+\frac{1}{2}} = (I \times J)^{t-1} \times I = H_t \times I ,$$

and denote by $H_{t+\frac{1}{2}}$ the finite field it generates. We have now defined H_s and H_s for all half-integers s , namely all^{7/} $s \in \mathbb{N}_2 \equiv \{1, 1\frac{1}{2}, 2, 2\frac{1}{2}, \dots\}$. Note that $\{H_s\}_{s \in \mathbb{N}_2}$ is an increasing sequence of finite subfields of H_∞ , converging to H_∞ as $s \rightarrow \infty$.

Since our probability space is actually $\Omega = H_\infty \times K$ and not H_∞ , we will denote the field generated by H_s on Ω also by H_s ; this will generate no confusion.

4.2 Banach Limit

In order to deal with the non-summability of the sequences of payoffs, we introduce the concept of a "Banach limit" (e.g., see Dunford and Schwartz [1958], p. 73).

As usual, let ℓ_∞ be the (Banach) space of all real bounded sequences $x = \{x_n\}_{n=1}^\infty$. A Banach limit is a real operator L on ℓ_∞ with the following properties^{8/} (holding for all $x = \{x_n\}_n$ and $y = \{y_n\}_n$ in ℓ_∞ , and λ, μ in \mathbb{R}):

$$(4.1) \quad L(\{\lambda \cdot x_n + \mu \cdot y_n\}) = \lambda \cdot L(\{x_n\}) + \mu \cdot L(\{y_n\}) ,$$

$$(4.2) \quad L(\{x_{n+1}\}_{n=1}^{\infty}) = L(\{x_n\}_{n=1}^{\infty}) ,$$

$$(4.3) \quad \liminf_{n \rightarrow \infty} x_n \leq L(\{x_n\}) \leq \limsup_{n \rightarrow \infty} x_n .$$

In particular, note that (4.3) implies:

$$(4.4) \quad L(\{x_n\}) = \lim_{n \rightarrow \infty} x_n, \text{ if } \{x_n\} \text{ is a convergent sequence} .$$

Therefore the Banach limit is an extension of the notion of limit (to all bounded sequences). To slightly simplify the notation, we will henceforth write $L[x_n]$ for $L(\{x_n\}_{n=1}^{\infty})$.

Three further properties of Banach limits will be needed.

Lemma 4.5: Let L be a Banach limit, and let $\{x_n\}_n, \{y_n\}_n \in \ell_{\infty}$. Then

$$|L[x_n] - L[y_n]| \leq \limsup_{n \rightarrow \infty} |x_n - y_n| .$$

Proof: Immediate by (4.1) and (4.3).

Q.E.D.

Lemma 4.6: Let L be a Banach limit, and $X = \{X_n\}_{n=1}^{\infty}$ an ℓ_{∞} -valued random variable (i.e., X is a measurable function from some probability space into ℓ_{∞}). If X has only finitely many values, then

$$L[E(X_n)] = E(L[X_n]) .$$

Proof: Immediate by (4.1).

Q.E.D.

In particular, this result will be useful for conditional expectations over finite fields. One could actually define a stronger concept of Banach limit, which commutes with the expectation operator for any uniformly bounded (or, even uniformly integrable) sequence of random variables--without the finiteness assumption. The construction of such a so-called "medial limit" requires however the use of the continuum hypothesis--and it is not needed in our proof (cf. Mokobodzki, see Meyer [1973]).

Lemma 4.7: Let L be a Banach limit, and C a compact and convex subset of some Euclidean space R^m . Let $\{x_n\}_{n=1}^{\infty}$ be a sequence in C , with $x_n = (\xi_n^{(1)}, \xi_n^{(2)}, \dots, \xi_n^{(m)})$. Let $\eta^{(r)} = L[\xi_n^{(r)}]$ for $r = 1, 2, \dots, m$. Then $y = (\eta^{(1)}, \eta^{(2)}, \dots, \eta^{(m)}) \in C$.

Proof: Let q be any vector in R^m , then by (4.1), (4.3) and $x_n \in C$,

$$q \cdot y = L[q \cdot x_n] \leq \limsup_{n \rightarrow \infty} q \cdot x_n \leq \sup \{q \cdot c : c \in C\}.$$

This holds for all q ; since C is a compact convex set, it implies $y \in C$.

Q.E.D.

Given a Banach limit L , we can now define the concept of an L-equilibrium point in $\Gamma_{\infty}(p)$, by replacing (2.7) with

$$(4.8) \quad L[E_{\sigma, \tau}^k(a_T^k)] \geq L[E_{\sigma', \tau}^k(a_T^k)] ,$$

and (2.8) with

$$(4.9) \quad L[E_{\sigma, \tau, p}(\beta_T)] \geq L[E_{\sigma', \tau, p}(\beta_T)] ,$$

where the limit L is taken with respect to the index $T = 1, 2, \dots$; this convention will be kept throughout this section. The corresponding payoffs will then be

$$(4.10) \quad L[E_{\sigma, \tau}^k(a_T^k)] = a^k$$

for each k in K , and

$$(4.11) \quad L[E_{\sigma, \tau, p}(\beta_T)] = \beta .$$

We put $a = (a^k)_{k \in K}$.

In view of (4.3), every equilibrium point is also an L -equilibrium point for any Banach limit L .

Throughout this section, we fix the following: a Banach limit L , a probability vector p in Δ^K , and an L -equilibrium point (σ, τ) in $\Gamma_{\infty}(p)$ with payoffs $(a, \beta) \in R_M^K \times R_M$. Unless stated otherwise, the probability measure $P \equiv P_{\sigma, \tau, p}$ is assumed (on the space Ω), with $E \equiv E_{\sigma, \tau, p}$ the corresponding expectation operator. Thus, all statements "a.s.", "martingale", and so on, will be with respect to P . Also, we will use E^k for the conditional expectation $E(\cdot | \kappa = k)$.

Our purpose is to construct a G-process starting at (a, β, p) . The probability space on which it will be defined is Ω , and the sequence of fields is $\{H_s\}_{s \in \mathbb{N}_2}$.

4.3 The Martingale $\{p_s\}$

For each $k \in K$, $s \in \mathbb{N}_2$ and an history $h_s \in H_s$, let $p_s^k = p_s^k(h_s)$ be the conditional probability of the "true" game κ of being k , given σ, τ, p and h_s (namely, if $s = t \in \mathbb{N}$, the first $t - 1$ moves of each player; if $s = t + \frac{1}{2}$, $t \in \mathbb{N}$, the first t moves of player 1 and $t - 1$ moves of player 2). We can thus write

$$p_s^k = P_{\sigma, \tau, p}(\kappa = k | H_s) = P(k | H_s)$$

(on each atom $h_s \in H_s$ of H_s , p_s^k is a.s. constant, thus a.s. equal to $p_s^k(h_s)$). We put $p_s = (p_s^k)_{k \in K}$.

Proposition 4.12: The sequence $\{p_s\}_{s \in \mathbb{N}_2}$ is a Δ^K -valued martingale with respect to $\{H_s\}_{s \in \mathbb{N}_2}$, satisfying:

$$(4.13) \quad p_1 \equiv p.$$

$$(4.14) \quad p_{t+\frac{1}{2}} = p_{t+1} \quad \text{for all } t \in \mathbb{N}.$$

(4.15) There exists a Δ^K -valued random variable p_∞ such that
 $p_s \rightarrow p_\infty$ a.s. as $s \rightarrow \infty$.

Proof: The fact that $\{p_s^k\}_s$ forms a martingale is immediate from its definition. Since it is bounded, it must converge a.s., say to p_∞^k ; then $p_\infty = (p_\infty^k)_{k \in K}$. (4.14) follows from the fact that given $h_{t+\frac{1}{2}}$ (actually, only h_t suffices), the t -th move j_t of player 2 is independent of κ ; as for (4.13)--at $t = 1$ there is no history yet, hence posteriors and priors coincide. Q.E.D.

4.4 The Martingales $\{\gamma_s\}$ and $\{\delta_s\}$

In Section 2, we defined the average payoffs of the two players up to time T (see (2.5) and (2.6)). We will find it useful to define also

$$(4.16) \quad \alpha_T = \frac{1}{T} \sum_{t=1}^T A^K(i_t, j_t)$$

(i.e., $\alpha_T = a_T^K$). For each $s \in N_2$, let

$$\gamma_s = L[E(\alpha_T | H_s)] ,$$

$$\delta_s = L[E(\beta_T | H_s)] .$$

Thus, γ_s and δ_s are the (Banach) limits of the expected average payoffs to player 1 and player 2, respectively, given a history h_s .

Proposition 4.17: The sequences $\{\gamma_s\}_{s \in \mathbb{N}_2}$ and $\{\delta_s\}_{s \in \mathbb{N}_2}$ are R_M -valued martingales with respect to $\{H_s\}_{s \in \mathbb{N}_2}$, satisfying:

$$(4.18) \quad \gamma_1 \equiv p \cdot a \quad \text{and} \quad \delta_1 \equiv \beta \quad .$$

$$(4.19) \quad \text{There exist } R_M\text{-valued random variables } \gamma_\infty \text{ and } \delta_\infty \text{ such that } \gamma_s \rightarrow \gamma_\infty \text{ and } \delta_s \rightarrow \delta_\infty \text{ a.s. as } s \rightarrow \infty \quad .$$

Proof: We can use Lemma 4.6--the field H_s being finite, γ_s has finitely many values:

$$\begin{aligned} E(\gamma_{s+1/2} | H_s) &= E(L[E(\alpha_T | H_{s+1/2})] | H_s) \\ &= L[E(E(\alpha_T | H_{s+1/2}) | H_s)] \\ &= L[E(\alpha_T | H_s)] = \gamma_s \quad . \end{aligned}$$

Thus $\{\gamma_s\}_{s \in \mathbb{N}_2}$ forms a martingale. It is bounded by M (which bounds all possible payoffs by (3.8), hence also averages, expectations and limits--by (4.3)--of those). Therefore it converges to some limit γ_∞ . For $s = 1$, we have

$$E(\alpha_T | H_1) = E(\alpha_T) = \sum_{k \in K} p^k E^k(a_T^k) \quad ,$$

hence (4.1) and (4.10) give $\gamma_1 \equiv p \cdot a$. The sequence $\{\delta_s\}$ is dealt with in a similar way ($\delta_1 \equiv \beta$ is just (4.11)). Q.E.D.

4.5 The Martingales $\{c_s\}$ and $\{d_s\}$

We now associate vector payoffs to each infinite history. We define, for each k in K and T in N ,

$$b_T^k = \frac{1}{T} \sum_{t=1}^T B^k(i_t, j_t) ,$$

in a similar way to the definition (2.5) of a_T^k . Note that these are random variables, H_{T+1} -measurable (k is fixed; in contrast, a_T and β_T in (4.16) and (2.6) are $(H_{T+1} \otimes 2^K)$ -measurable). We further remark that a_T^k and b_T^k are defined for all histories--even those which may be incompatible with $\kappa = k$ according to (σ, τ) .

If the limit of a_T^k (as $T \rightarrow \infty$) would always exist, it would imply $E(\lim a_T^k) = \lim E(a_T^k)$. However, this is not the case, and the Banach limit L commutes with the expectation operator if there are only finitely many values (see Lemma 4.6 and the discussion thereafter). We define, for each $s \in N_2$:

$$c_s^k = L[E(a_T^k | H_s)]$$

$$d_s^k = L[E(b_T^k | H_s)]$$

(again, for each k in K). Note that the expectations are not conditional on $\kappa = k$; thus, the probability of any history is its total probability, summed over all k in K .

One can interpret c_s^k and d_s^k as follows. Let $h_s \in H_s$ have positive probability, then $p_s = p_s(h_s)$ is the vector of posterior probabilities for the various games k . Assume that after h_s occurred, player 1 replaces his strategy σ by his average non-revealing strategy there; namely, for all k in K , he uses $\sum_{k \in K} p_s^k \sigma_t(h_t; k)$ instead of $\sigma_t(h_t; k)$ whenever $t \geq s$ and h_t coincides with h_s up to s . The expected average payoffs up to T in game k will then be $E(a_T^k | h_s)$ and $E(b_T^k | h_s)$, respectively. As we shall see later, the difference in payoffs due to this change in strategy becomes negligible as $s \rightarrow \infty$ (Proposition 4.23). Intuitively, this is due to the fact that after sufficiently many stages, player 1 has already revealed (almost) everything he is ever going to reveal about the true game κ ; therefore, he must thereafter play (almost) non-revealing, or (almost) independent of κ . In technical terms, this occurs whenever the martingales are close to their limits.

As usual, we write c_s for $(c_s^k)_{k \in K}$ and d_s for $(d_s^k)_{k \in K}$. The set F was defined in (3.7) as the set of all "feasible" vector payoffs to both players (in the one-shot game).

Proposition 4.20: The sequences $\{c_s\}_{s \in \mathbb{N}_2}$ and $\{d_s\}_{s \in \mathbb{N}_2}$ are r_M^K -valued martingales with respect to $\{H_s\}_{s \in \mathbb{N}_2}$, satisfying:

(4.21) There exist R_M^K -valued random variables c_∞ and d_∞ such that $c_s \rightarrow c_\infty$ and $d_s \rightarrow d_\infty$ a.s. as $s \rightarrow \infty$.

(4.22) $(c_\infty, d_\infty) \in F$ a.s. .

Proof: The martingale property and (4.21) are proved in a similar way to Proposition 4.17. For every T in N , the vector $((a_T^k)_{k \in K}, (b_T^k)_{k \in K})$ belongs to the compact convex set F , as an average of such vectors. The same holds for its expectations, and by Lemma 4.7 for its Banach limits (c_s, d_s) too. (4.22) now follows by letting $s \rightarrow \infty$. Q.E.D.

The next proposition makes precise the statement that, as $s \rightarrow \infty$, player 1 plays "almost" non-revealing after s (see the discussion following the definition of c_s^k and d_s^k).

Proposition 4.23:

$$\gamma_s - p_s \cdot c_s \rightarrow 0$$

$$\delta_s - p_s \cdot d_s \rightarrow 0$$

a.s. as $s \rightarrow \infty$.

Proof: We prove here the first part. Fix $s \in N_2$, and let $t \geq s$, $t \in N$. Conditioning over H_{t+1} and κ gives (recall that $p_{t+1}^k = P(\kappa = k | H_{t+1})$):

$$\begin{aligned}
 E(A^k(i_t, j_t) | H_s) &= E\left(\sum_{k \in K} p_{t+1}^k A^k(i_t, j_t) | H_s\right) \\
 &= \sum_{k \in K} p_s^k E(A^k(i_t, j_t) | H_s) \\
 &\quad + \sum_{k \in K} E((p_{t+1}^k - p_s^k) A^k(i_t, j_t) | H_s) .
 \end{aligned}$$

We sum this for all t in the range $s \leq t \leq T$, and note that total payoffs up to s are bounded by sM , to obtain

$$\begin{aligned}
 &|E(\alpha_T | H_s) - \sum_{k \in K} p_s^k E(a_T^k | H_s)| \\
 &\leq 2\frac{sM}{T} + \frac{1}{T}M \sum_{\substack{t \in N \\ s \leq t \leq T}} \sum_{k \in K} E(|p_{t+1}^k - p_s^k| | H_s) .
 \end{aligned}$$

We denote (for each $s \in N_2$)

$$\pi_s = \sum_{k \in K} \sup_{\substack{t \in N \\ t \geq s}} |p_{t+1}^k - p_s^k| ,$$

and let $T \rightarrow \infty$. By Lemma 4.5 and (4.1),

$$|\gamma_s - p_s \cdot c_s| \leq M \cdot E(\pi_s | H_s) .$$

Since $\{p_s\}_s$ converges a.s. as $s \rightarrow \infty$ by (4.15), it follows that $E(\pi_s | H_s) \rightarrow 0$ a.s. as $s \rightarrow \infty$; this assertion is proved in the next lemma.

Q.E.D.

Lemma 4.24: Let $\{X_n\}_{n=1}^{\infty}$ be a bounded sequence of real random variables, converging a.s. as $n \rightarrow \infty$, and let $\{F_n\}_{n=1}^{\infty}$ be a non-decreasing sequence of σ -fields. Define

$$Y_n = \sup_{m \geq n} |X_m - X_n|,$$

then

$$E(Y_n | F_n) \rightarrow 0$$

a.s. as $n \rightarrow \infty$.

Proof: Let $X_{\infty} = \lim X_n$, and put $Z_n = \sup_{m \geq n} |X_m - X_{\infty}|$. Then $\{Z_n\}_{n=1}^{\infty}$ is a non-increasing sequence as $n \rightarrow \infty$, converging a.s. to zero. Therefore $\{E(Z_n | F_n)\}_{n=1}^{\infty}$ is a bounded super-martingale^{2/} with respect to $\{F_n\}_{n=1}^{\infty}$, hence converges a.s. to some Z_{∞} . Now $E(E(Z_n | F_n)) = E(Z_n) \rightarrow 0$, thus $E(Z_{\infty}) = 0$ and $Z_{\infty} = 0$ a.s. Noting that $Y_n \leq 2Z_n$ completes the proof. Q.E.D.

Finally we have

Corollary 4.25:

$$\gamma_{\infty} = p_{\infty} \cdot c_{\infty} \text{ a.s.}$$

$$\delta_{\infty} = p_{\infty} \cdot d_{\infty} \text{ a.s.}$$

Proof: (4.15), (4.19), (4.21) and Proposition 4.23. Q.E.D.

4.6 The Martingales $\{e_s\}$ and $\{f_s\}$

For each k in K and s in N_2 , we define

$$e_s^k = \sup_{\sigma'} L[E_{\sigma', \tau}^k(a_T^k | H_s)] ,$$

where σ' ranges over all strategies of player 1 (note that the expectation now is conditional on $\kappa = k$). Thus, for every history $h_s \in H_s$, e_s^k is the most player 1 can obtain if the true game is k and player 2 uses τ --given that h_s has already occurred.

Proposition 4.26: For every k in K , s in N_2 and t in N :

$$(4.27) \quad e_1^k \equiv a^k .$$

$$(4.28) \quad e_s^k \geq c_s^k .$$

$$(4.29) \quad e_{t+1/2}^k = E(e_{t+1}^k | H_{t+1/2}) .$$

$$(4.30) \quad e_t^k(h_t) = \max_{i_t \in I} e_{t+1/2}^k(h_t, i_t) \text{ for all } h_t \text{ in } H_t .$$

Proof: (4.27) is just (4.8) and (4.10). To obtain (4.28), we consider the following σ' : if $\kappa = k$ and h_s occurred, play the average non-revealing strategy given by σ ; namely,

$\sigma'_t(h_t; k) = \sum_{k' \in K} p_s^{k'} \sigma_t(h_t; k')$ for all $t \geq s$ and h_t in H_t that coincide with h_s up to s (see the discussion following the definition of c_s^k and d_s^k in subsection 4.5).

To prove (4.29), note that the additional information from $t + \frac{1}{2}$ to $t + 1$ is j_t -- whose distribution depends on τ and h_t only, hence is the same in $E_{\sigma', \tau}^k$ as in E . Therefore

$$\begin{aligned} & \sup_{\sigma'} L[E_{\sigma', \tau}^k(a_T^k | h_{t+\frac{1}{2}})] \\ &= \sup_{\sigma'} L[E(E_{\sigma', \tau}^k(a_T^k | h_{t+\frac{1}{2}}, j_t) | h_{t+\frac{1}{2}})] \\ &= \sup_{\sigma'} E(L[E_{\sigma', \tau}^k(a_T^k | h_{t+1})] | h_{t+\frac{1}{2}}) \end{aligned}$$

(we used Lemma 4.6). Given $h_{t+\frac{1}{2}}$, the first stage player 1 has to choose a move in $t + 1$, and by that time he will already know j_t . Thus, the best he can do given $h_{t+\frac{1}{2}}$ is just to do his best given $(h_{t+\frac{1}{2}}, j_t) = h_{t+1}$, for each possible j_t . Therefore, the last expression is

$$= E(\sup_{\sigma'} L[E_{\sigma', \tau}^k(a_T^k | h_{t+1})] | h_{t+\frac{1}{2}}),$$

proving (4.29).

Next, let $h_t \in H_t$ be given. For any σ' , its relevant part for $L[E_{\sigma', \tau}^k(a_T^k | h_t)]$ consists of a probability distribution $\pi = \sigma'_t(h_t; k)$ in Δ^I for choosing i_t , and some strategy afterwards

$\sigma'' \equiv \sigma''(i_t) = \sigma'((h_t, i_t, \cdot); k)$, for each possible i_t . Therefore (again using Lemma 4.6)

$$e_t^k(h_t) = \sup_{\pi \in \Delta} \sup_{i_t \in I} E^{k, \pi}(L[E^{k, \pi}(a_t^k | h_t, i_t)] | h_t) ,$$

where $E^{k, \pi}$ is just $E_{\sigma', \tau}^k$. The choice of $\sigma''(i_t)$ can be done separately for each i_t , therefore we can interchange the first $E^{k, \pi}$ with the supremum over $\sigma''(i_t)$, to obtain

$$e_t^k(h_t) = \sup_{\pi \in \Delta} \sum_{i_t \in I} \pi(i_t) e_{t+1/2}^k(h_t, i_t) .$$

The supremum is attained by giving positive probability $\pi(i_t)$ only to those i_t for which $e_{t+1/2}^k(h_t, i_t)$ is maximal; this proves (4.30).

Q.E.D.

It is easy to see that (4.29) and (4.30) imply that $e_s = (e_s^k)_{k \in K}$ forms a super-martingale. To obtain a martingale, we define

$$f_s^k = e_s^k + \sum_{\substack{r \in N \\ r < s}} (e_r^k - e_{r+1/2}^k)$$

for all k in K and s in N_2 , and put $f_s = (f_s^k)_{k \in K}$.

Proposition 4.31: The sequence $\{f_s\}_{s \in N_2}$ is an R_M^K -valued martingale with respect to $\{H_s\}_{s \in N_2}$, satisfying:

$$(4.32) \quad f_1 \equiv a .$$

$$(4.33) \quad f_t = f_{t+1/2} \text{ for all } t \text{ in } N .$$

$$(4.34) \quad \text{There exists an } R_M^K\text{-valued random variable } f_\infty \text{ such that } f_s \rightarrow f_\infty \text{ a.s. as } s \rightarrow \infty .$$

$$(4.35) \quad f_s \geq e_s \geq c_s \text{ for all } s \text{ in } N_2 .$$

$$(4.36) \quad f_\infty \geq c_\infty \text{ and } p_\infty \cdot f_\infty = p_\infty \cdot c_\infty \text{ a.s. .}$$

Proof: (4.32) is immediate from (4.27) and the definition of f_1 . Let $t \in N$, then

$$\begin{aligned} f_{t+1/2}^k &= e_{t+1/2}^k + \sum_{\substack{r \in N \\ r < t}} (e_r^k - e_{r+1/2}^k) + (e_t^k - e_{t+1/2}^k) \\ &= e_t^k + \sum_{\substack{r \in N \\ r < t}} (e_r^k - e_{r+1/2}^k) = f_t^k , \end{aligned}$$

proving (4.33). Moreover,

$$f_{t+1}^k - f_{t+1/2}^k = e_{t+1}^k - e_{t+1/2}^k ,$$

since the sum of the additional terms is identical in both f_{t+1}^k and $f_{t+\frac{1}{2}}^k$. Using (4.29) completes the proof that $\{f_s\}$ indeed forms a martingale (note: with respect to the probability measure $P \equiv P_{\sigma, \tau, p}$). Since it is bounded, (4.34) follows.

By (4.30), $f_s^k \geq e_s^k$ (all the additional terms are non-negative), hence $f_s^k \geq e_s^k \geq c_s^k$ and $f_\infty^k \geq c_\infty^k$ a.s. (recall (4.28) and (4.21)).

Therefore $p_\infty \cdot f_\infty \geq p_\infty \cdot c_\infty = \gamma_\infty$ (by Corollary 4.25). To obtain the opposite inequality, we note that $\{p_s\}$ and $\{f_s\}$ form a bi-martingale, hence $\{p_s \cdot f_s\}$ is a martingale (Proposition 3.18), and we have by (4.13), (4.32) and (4.18)

$$E(p_\infty \cdot f_\infty) = E(p_1 \cdot f_1) = p \cdot a = E(\gamma_1) = E(\gamma_\infty) ,$$

proving that $p_\infty \cdot f_\infty = \gamma_\infty = p_\infty \cdot c_\infty$ a.s. .

Q.E.D.

4.7 Individual Rationality

We start with player 1.

Proposition 4.37: For all vectors q in Δ^K and all s in N_2

$$q \cdot f_s \geq q \cdot e_s \geq (\text{val}_1 A)(q) .$$

Proof: Let $q \in \Delta^K$, and consider the one-shot zero-sum game $A(q)$. By definition (3.1), player 1 has a strategy $u \in \Delta^I$ such that for any strategy $v \in \Delta^J$ of player 2,

$$(4.38) \quad (\text{val}_1 A)(q) \leq \sum_{i \in I} \sum_{j \in J} u_i v_j \sum_{k \in K} q^k A^k(i, j) .$$

Let $h_s \in H_s$ have positive probability (under (σ, τ)). Define a new strategy σ' of player 1 as follows: $\sigma'((h_s, \cdot); k) = u$ for all k , and σ' equals σ otherwise (thus, after h_s has occurred, player 1 makes independent randomizations with distribution u at all stages and all k). By (4.38), we have for all t in N , $t \geq s$

$$(\text{val}_1 A)(q) \leq E_{\sigma', \tau} \left(\sum_{k \in K} q^k A^k(i_t, j_t) \mid h_s \right) .$$

As $T \rightarrow \infty$, payoffs before s become negligible in a_T^k , and we have (by (4.2), (4.3) and then (4.1)):

$$\begin{aligned} (\text{val}_1 A)(q) &\leq L[E_{\sigma', \tau} \left(\sum_{k \in K} q^k a_T^k \mid h_s \right)] \\ &= \sum_{k \in K} q^k L[E_{\sigma', \tau} (a_T^k \mid h_s)] . \end{aligned}$$

Recalling the definition of e_s^k (note that given h_s , $E_{\sigma', \tau}^k$ is independent of k),

$$(\text{val}_1 A)(q) \leq \sum_{k \in K} q^k e_s^k(h_s) = q \cdot e_s(h_s) ,$$

and $q \cdot e_s \leq q \cdot f_s$ follows from (4.35) .

Q.E.D.

Corollary 4.39: For all q in Δ^K

$$q \cdot f_\infty \geq (\text{val}_1 A)(q) \text{ a.s. .}$$

Proof: (4.34) and Propositions 3.16(i) and 4.37. Q.E.D.

We consider now player 2.

Proposition 4.40: For all s in N_2

$$\delta_s \geq (\text{vex val}_2 B)(p_s) \text{ a.s. .}$$

Proof: For each q in Δ^K , let $\Gamma_\infty^2(q)$ be defined in the same way as $\Gamma_\infty(q)$, but with payoff matrices $(-B^k)_{k \in K}$ instead of $(A^k)_{k \in K}$ for player 1. This is a zero-sum repeated game; therefore player 2 (the uninformed player) has a strategy $\hat{\tau} \equiv \hat{\tau}(q)$ such that

$$\begin{aligned} \limsup_{T \rightarrow \infty} E_{\sigma', \hat{\tau}, q} \left(\frac{1}{T} \sum_{t=1}^T (-B^k(i_t, j_t)) \right) &\leq (\text{cav} (\text{val}_1 (-B)))(q) \\ &= -(\text{vex val}_2 B)(q) \end{aligned}$$

for all σ' (cf. Aumann and Maschler [1966] -- $\hat{\tau}$ may be taken to be the corresponding Blackwell strategy; "cav" is the concavification of a function, and we use (3.2)). Thus,

$$(4.41) \quad \liminf_{T \rightarrow \infty} E_{\sigma', \hat{\tau}, q} (\beta_T) \geq (\text{vex val}_2 B)(q) \text{ .}$$

Let $h_s \in H_s$ have positive probability under (σ, τ) , and consider the following strategy τ' of player 2: after h_s has occurred, τ' is

$\hat{\tau}(p_s)$, where $p_s = p_s(h_s)$ is the vector of posterior probabilities given h_s ; otherwise, τ' equals τ . Let $T \geq s$, then we condition on H_s , to obtain

$$E(\beta_T) - E'(\beta_T) = P(h_s)(E(\beta_T|h_s) - E'(\beta_T|h_s)) ,$$

where $E' = E_{\sigma, \tau', p}$ (up to stage s -- no difference between E and E' ; afterwards--only if h_s has occurred). Apply the Banach limit L as $T \rightarrow \infty$; since (σ, τ) is an equilibrium (see (4.9)), we get by (4.1)

$$0 \leq P(h_s)(L[E(\beta_T|h_s)] - L[E'(\beta_T|h_s)]) ,$$

hence

$$\delta_s(h_s) = L[E(\beta_T|h_s)] \geq L[E'(\beta_T|h_s)]$$

(since $P(h_s) > 0$). By (4.3) and (4.41) (with $\sigma' = \sigma$; note that payoffs up to s do not matter as $T \rightarrow \infty$), the proof is completed. Q.E.D.

Corollary 4.42: $\delta_\infty \geq (\text{vex val}_2 B)(p_\infty)$ a.s. .

Proof: (4.15), (4.19) and Propositions 3.16(ii) and 4.40. Q.E.D.

4.8 The G-Process

We have thus completed the proof of

Proposition 4.43: Let (a, β) be the payoffs of an L-equilibrium point (σ, τ) in $\Gamma_\infty(p)$. Then there exists a G-process starting at (a, β, p) .

Proof: The probability space is $(\Omega, H_\infty, P_{\sigma, \tau, p})$; the sequence of fields is $\{H_s\}_{s \in \mathbb{N}_2}$, and the G-process $\{g_s\}_{s \in \mathbb{N}_2}$ is given by $g_s = (f_s, \delta_s, p_s)$. All the required properties are indeed satisfied: $g_1 \equiv (a, \beta, p)$ by (4.13), (4.18) and (4.32); the limit $g_\infty = (f_\infty, \delta_\infty, p_\infty)$ (see (4.15), (4.19) and (4.34)) belongs to the set G a.s. by (4.22), (4.36) and Corollaries 4.25, 4.39 and 4.42; and finally the "bi" property (3.15) is given in (4.14) and (4.33). Q.E.D.

5. From Martingale to Equilibrium

This section is devoted to the proof of the second half of our result; namely, given a G-process the corresponding uniform equilibrium point is constructed.

Let $g = (a, \beta, p)$ belong to G^* . Thus, we are given a probability space^{10/} (Z, \mathcal{Z}, Q) , a non-decreasing sequence $\{Z_n\}_{n=1}^\infty$ of finite subfields of Z , and a G-process $\{g_n\}_{n=1}^\infty = \{(f_n, \delta_n, p_n)\}_{n=1}^\infty$ with respect to $\{Z_n\}_{n=1}^\infty$, starting at g ; i.e.,

$$(5.1) \quad (f_1, \delta_1, p_1) \equiv (a, \beta, p) \quad Q - \text{a.s.}$$

Without loss of generality, we will assume that Z_1 is the trivial field $\{Z, \phi\}$. Let $g_\infty = (f_\infty, \delta_\infty, p_\infty)$ be a Q -a.s. limit of g_n as $n \rightarrow \infty$; then $g_\infty \in G$ a.s. . We will find it useful to weaken the

bi-property (3.15) to the following:

$$(5.2) \quad \|f_{n+1} - f_n\| \cdot \|p_{n+1} - p_n\| = 0 \quad \text{a.s. for all } n = 1, 2, \dots$$

This means that on each atom of Z_n , either f_{n+1} is constant (and thus equals f_n), or p_{n+1} is constant (and equals p_n); however, which one of the two is true may differ from one atom to the other. It is easy to see that G^* does not change (to obtain (3.15) from (5.2), insert between each Z_n and Z_{n+1} an additional field $Z_{n+\frac{1}{2}}$, and put $g_{n+\frac{1}{2}} = g_{n+1}$ if $f_{n+1} = f_n$ and $g_{n+\frac{1}{2}} = g_n$ otherwise).

5.1 Standard G-Process

To simplify the construction of the equilibrium point, we will work with a G-process having the following additional property:

$$(5.3) \quad \begin{aligned} &\text{For every atom } z_n \text{ of } Z_n \text{ there are exactly two} \\ &\text{atoms } z'_{n+1} \text{ and } z''_{n+1} \text{ of } Z_{n+1} \text{ contained in } z_n, \\ &\text{and } Q(z'_{n+1} | z_n) = Q(z''_{n+1} | z_n) = 1/2. \end{aligned}$$

Such a G-process will be called standard.

Proposition 5.4: For every g in G^* there exists a standard G-process starting at g .

Proof: We will show how to "transform" any G-process into a standard one.

Given a G-process $\{g_n\}_{n=1}^{\infty}$ with respect to $\{Z_n\}_{n=1}^{\infty}$, we can describe the sequence of fields as a "probability tree" as follows. The nodes in the n-th layer are the atoms of Z_n ; the root (i.e., the first layer) can be taken to be Z (by (5.1)). A (directed) arc leads from an atom z_n of Z_n to an atom z_m of Z_m if and only if $m = n + 1$ and $z_m \equiv z_{n+1} \subset z_n$. We associate the probability $Q(z_{n+1}|z_n)$ to this arc and define the probability of a finite path starting at the root to be the product of the probabilities of all its arcs. This clearly equals $Q(z_n)$, where z_n is the endpoint of the path. This probability distribution is then uniquely extended in a standard way to all infinite paths in the tree starting at the root; we will denote this probability measure also by Q . This completes the description of our probability tree.

The G-process $\{g_n\}_{n=1}^{\infty}$ can now be regarded as being defined on the nodes of the tree; we will write $g_n(z_n)$ for the value of g_n on the atom z_n of Z_n . The properties (3.12)-(3.14) and (5.2) defining a G-process become:

- (i) $g_1(z_1) = g$.
- (ii) $E(g_{n+1}(z_{n+1})|z_{n+1} \text{ succeeds } z_n) = g_n(z_n)$ for all z_n .
- (iii) The sequence $\{g_n(z_n)\}_{n=1}^{\infty}$ converges for almost all infinite paths, and the limit g_{∞} belongs to G a.s. .
- (iv) For each node z_n , either $f_{n+1}(z_{n+1}) = f_n(z_n)$ for all successors z_{n+1} of z_n , or $p_{n+1}(z_{n+1}) = p_n(z_n)$ for all successors z_{n+1} of z_n .

In order to obtain property (5.3), we need two kinds of modifications of the tree--and thus, of the G-process. First, we make the number of (immediate) successors of each node exactly two; and second, we make the probability of every arc precisely $1/2$.

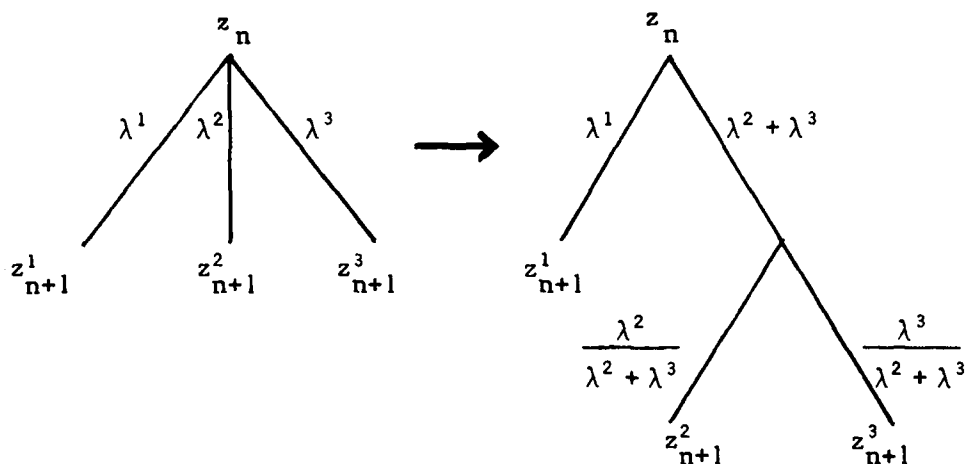
For the former, we have two cases. If there is only one successor z_{n+1} of z_n , we can add an additional copy of the whole subtree starting at z_{n+1} , and thus obtain two successors z_{n+1} and z'_{n+1} (which is identical to z_{n+1}) --and moreover with probability $1/2$ each (from z_n). Now, assume z_n has more than two successors, say $\{z_{n+1}^r\}_{r=1}^m$. We then introduce additional nodes in between; e.g., at level $n+1$ we will have z_{n+1}^1 and the union of $z_{n+1}^2, \dots, z_{n+1}^m$; from the latter, at level $n+2$ we will have z_{n+1}^2 and the union of $z_{n+1}^3, \dots, z_{n+1}^m$; and so on. The probabilities of the new arcs will be defined as the corresponding conditional probabilities; the value of the G-process at the new nodes, as the conditional expectation. As an example, see Figure 2; the value of the G-process at the new node will be

$$\frac{\lambda^2}{\lambda^2 + \lambda^3} g_{n+1}(z_{n+1}^2) + \frac{\lambda^3}{\lambda^2 + \lambda^3} g_{n+1}(z_{n+1}^3) .$$

Clearly, all four properties (i)-(iv) continue to hold after such modifications.

Next we have to make the probabilities of all arcs precisely $1/2$. Let z'_{n+1} and z''_{n+1} be the two successors of z_n , and let λ'

Figure 2



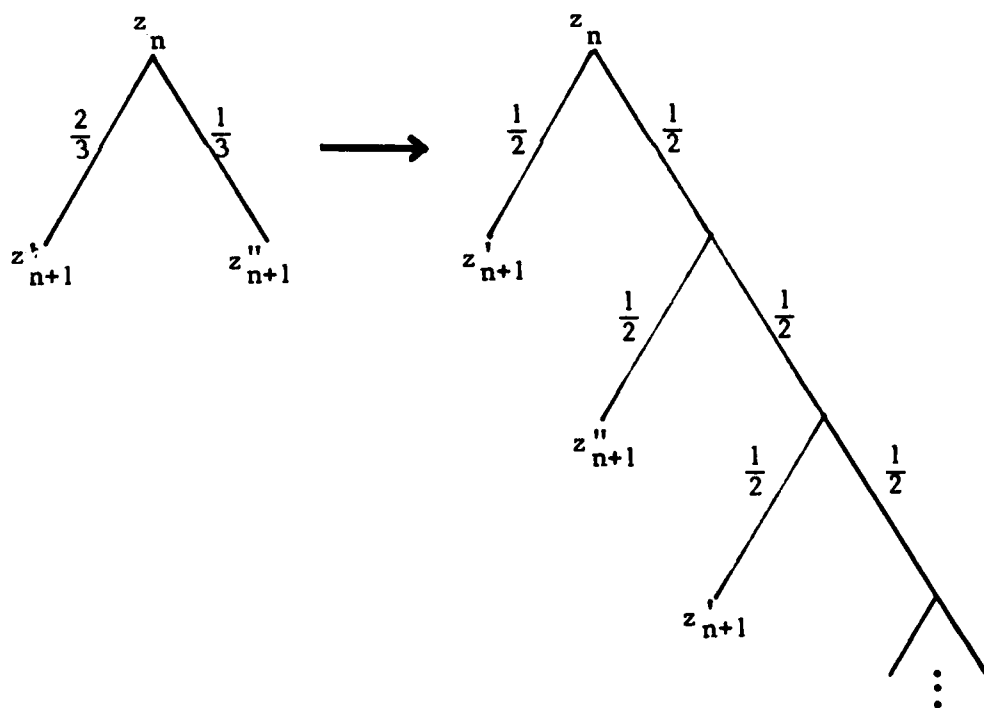
and $\lambda'' = 1 - \lambda'$ be the corresponding probabilities. We want to obtain z'_{n+1} with probability λ' and z''_{n+1} with probability λ'' by using the probability $1/2$ only. This is done as follows: we express λ' as a binary fraction

$$\lambda' = \sum_{m=1}^{\infty} \frac{1}{2^m} \lambda_m,$$

with $\lambda_m = 0$ or 1 for all m . We then consider an infinite sequence of independent Bernoulli trials, with "success" and "failure" having probability $1/2$ each, up to the first occurrence of "success". If this happens after m trials, then z'_{n+1} "results" if $\lambda_m = 1$ and z''_{n+1} "results" if $\lambda_m = 0$. Thus, the total probability of z'_{n+1} is precisely λ' (since the first "success" occurs at the m -th trial with probability

$1/2^m$), and that of z''_{n+1} is λ'' . This structure now replaces the original randomization between z'_{n+1} and z''_{n+1} in the tree. As an example, see Figure 3 (note that $\lambda' = 2/3$ gives $\lambda_m = 1$ for m odd and $\lambda_m = 0$ for m even). Again, the value of the G-process at a new node is the corresponding expectation.

Figure 3



With probability one, either z'_{n+1} or z''_{n+1} will be reached. If we do this modification at all nodes, the properties (i)-(iv) will not be affected (there are only countable many nodes, hence the probability of "success" not occurring in even case is still zero). Q.E.D.

Henceforth we will assume that the G-process we start with is already standard.

5.2 The Sequence $\{\theta_n\}$

The limit g_∞ of the G-process belongs to G a.s.; by (3.9), a corresponding point in F is thus obtained--and from it, a point in the set $\Delta^{I \times J}$ of "feasible joint actions".

For $\theta = (\theta(i,j))_{i \in I, j \in J}$ in $\Delta^{I \times J}$ and k in K , we will denote

$$A^k(\theta) = \sum_{i \in I} \sum_{j \in J} \theta(i,j) A^k(i,j) ,$$

and $A(\theta) = (A^k(\theta))_{k \in K}$; similarly for B .

Proposition 5.5: There exists a $\Delta^{I \times J}$ -valued random variable θ_∞ satisfying Q - a.s.

$$(5.6) \quad f_\infty \geq A(\theta_\infty) \quad \text{and} \quad p_\infty \cdot f_\infty = p_\infty \cdot A(\theta_\infty) ,$$

$$(5.7) \quad \delta_\infty = p_\infty \cdot B(\theta_\infty) .$$

Proof: By definition, $g_\infty \in G$ implies the existence of $(c_\infty, d_\infty) \in F$ satisfying $f_\infty \geq c_\infty$, $p_\infty \cdot f_\infty = p_\infty \cdot c_\infty$ and $\delta_\infty = p_\infty \cdot d_\infty$. Since F is precisely the set of $(A(\theta), B(\theta))$ for all θ in $\Delta^{I \times J}$, there is θ_∞ in $\Delta^{I \times J}$ such that $c_\infty = A(\theta_\infty)$ and $d_\infty = B(\theta_\infty)$. The measurability is obtained by the Measurable Selection Theorem (e.g., see Hildenbrand [1974]). Q.E.D.

Proposition 5.8: There exists a sequence $\{\theta_n\}_{n=1}^\infty$ of $\Delta^{I \times J}$ -valued random variables, satisfying Q - a.s. for all n in N and (i, j) in $I \times J$:

$$(5.9) \quad \theta_n \text{ is } Z_n\text{-measurable}.$$

$$(5.10) \quad |\theta_n(i, j) - E(\theta_\infty(i, j) | Z_n)| \leq \frac{1}{n}.$$

$$(5.11) \quad \theta_n \rightarrow \theta_\infty \text{ as } n \rightarrow \infty.$$

$$(5.12) \quad n\theta_n(i, j) \text{ is an integer}.$$

Proof: Define $\bar{\theta}_n = E(\theta_\infty | Z_n)$, then $\{\bar{\theta}_n\}_{n=1}^\infty$ forms a martingale converging to θ_∞ . Choose θ_n to be a rational approximation to $\bar{\theta}_n$ with denominators n (e.g., let $\theta'_n(i, j) = [n\bar{\theta}_n(i, j)]/n$, where $[x]$ denotes the largest integer not exceeding x , then $\theta_n(i, j)$ is either $\theta'_n(i, j)$ or $\theta'_n(i, j) + 1/n$, so as to have the sum equal 1). Q.E.D.

5.3 The Strategies σ and τ

We can now define the pair of strategies (σ, τ) . In a similar way to the so-called "Folk Theorem" for repeated games with complete

information (for a detailed proof, see the Lecture Notes of Hart [1980, Section IV]), they are based on a master plan and punishments. Each player follows the master plan as long as the other one does it too (at least, as long as no deviation is detected), and uses the corresponding punishment otherwise.

The master plan consists of two parts. Stages $t = n!$, for all $n = 1, 2, \dots$, are communication stages; the moves made serve as a mean of transmitting information (from the informed to the uninformed player), or of making a joint decision. All the other stages are payoff periods; well determined moves (namely, pure) are used in order for both players to accumulate the "right" payoffs. The sequence $n!$ was chosen since $(n-1)!$ is negligible relative to $n!$ as n goes to infinity--thus only the last period^{11/} really counts. Any other sequence with the same property could be used just as well.

The master plan is derived from the G-process. The moves at stage $t = n!$ correspond to the arcs from z_n to z_{n+1} in the tree (see subsection 5.1), whereas at stages $(n-1)! < t < n!$, one "stays" at z_n . Thus, a function ζ is defined inductively from the set of finite histories in the game for which no deviations occurred, to the set of atoms of the fields $\{Z_n\}$ --or, equivalently, to the set of nodes in the tree.

Let $i' \neq i''$ be two elements of I , and $j' \neq j''$ two elements of J , fixed throughout the remainder of this section. These two (pure) moves for each player will actually be their communication alphabet (thus, they essentially "talk" in a binary language).^{12/}

Let $t = n!$, let h_t be a history with no deviation from the master plan, and let $z_n = \zeta(h_t)$ be the corresponding node in the tree. We will define now the behaviour of each player at stage t , and also the resulting $\zeta(h_{t+1})$. If $p_{n+1} = p_n$ and $f_{n+1} = f_n$ at both nodes z'_{n+1} and z''_{n+1} succeeding z_n , the two players play arbitrarily at $t = n!$, and for all i_t and j_t , $\zeta(h_{t+1}) = z'_{n+1}$, say. Otherwise, we distinguish two cases:

- (i) $p_{n+1} \neq p_n$ (and thus $f_{n+1} = f_n$),
- (ii) $f_{n+1} \neq f_n$ (and thus $p_{n+1} = p_n$).

In case (i), we define for each k in K

$$\sigma(h_t; k)(i_t) = \begin{cases} \frac{p_{n+1}^k(z'_{n+1})}{2p_n^k(z_n)} & , \text{ if } i_t = i' , \\ \frac{p_{n+1}^k(z''_{n+1})}{2p_n^k(z_n)} & , \text{ if } i_t = i'' , \\ 0 & , \text{ otherwise .} \end{cases}$$

Since $p_n(z_n) = (p_{n+1}(z'_{n+1}) + p_{n+1}(z''_{n+1}))/2$ by (5.3), $\sigma(h_t; k)$ is indeed a probability distribution over I . As for player 2, we let $\tau(h_t)$ be arbitrary in this case, and then for all j_t in J , we put $\zeta(h_t, i', j_t) = z'_{n+1}$ and $\zeta(h_t, i'', j_t) = z''_{n+1}$.

Lemma 5.13: Assume that $P(\kappa = k | h_t) = p_n^k(\zeta(h_t))$ for all $k \in K$.

Then

$$P(\zeta(h_{t+1}) = z'_{n+1} | h_t) = P(\zeta(h_{t+1}) = z''_{n+1} | h_t) = \frac{1}{2}$$

and

$$P(\kappa = k | h_{t+1}) = p_n^k(\zeta(h_{t+1})).$$

Proof: Assume $i_t = i'$, then we have

$$\begin{aligned} P(i_t = i' | h_t) &= \sum_{k \in K} P(i_t = i' | h_t, \kappa = k) \cdot P(\kappa = k | h_t) \\ &= \sum_{k \in K} \frac{p_{n+1}^k(z'_{n+1})}{2p_n^k(z_n)} \cdot p_n^k(z_n) \\ &= \frac{1}{2} \sum_{k \in K} p_{n+1}^k(z'_{n+1}) = \frac{1}{2}. \end{aligned}$$

Therefore

$$\begin{aligned} P(\kappa = k | h_{t+1}) &= P(\kappa = k | h_t, i_t = i') \\ &= \frac{P(i_t = i' | h_t, \kappa = k) \cdot P(\kappa = k | h_t)}{P(i_t = i' | h_t)} \\ &= \frac{\frac{p_{n+1}^k(z'_{n+1})}{2p_n^k(z_n)} \cdot p_n^k(z_n)}{\frac{1}{2}} = p_{n+1}^k(z'_{n+1}). \end{aligned}$$

Similarly for $i_t = i''$.

Q.E.D.

Thus, if the posterior probabilities for the various values of k at stage $t = n!$ are p_n^k , then the new posteriors generated after the moves at time t are precisely p_{n+1}^k . Case (i) therefore corresponds to a transmission of information (about the value k of κ) from player 1 to player 2; we will henceforth call this signalling (by player 1).

In case (ii), $f_{n+1} \neq f_n$ and $p_{n+1} = p_n$; we define for all k in K

$$\sigma(h_t; k)(i_t) = \begin{cases} \frac{1}{2} & , \text{ if } i_t = i' & , \\ \frac{1}{2} & , \text{ if } i_t = i'' & , \\ 0 & , \text{ otherwise } & , \end{cases}$$

$$\tau(h_t)(j_t) = \begin{cases} \frac{1}{2} & , \text{ if } j_t = j' & , \\ \frac{1}{2} & , \text{ if } j_t = j'' & , \\ 0 & , \text{ otherwise } & , \end{cases}$$

and then $\zeta(h_t, i', j') = \zeta(h_t, i'', j'') = z'_{n+1}$, $\zeta(h_t, i', j'') = \zeta(h_t, i'', j') = z''_{n+1}$.

Lemma 5.14:

$$\begin{aligned} P(\zeta(h_{t+1}) = z_{n+1} | h_t) &= P(\zeta(h_{t+1}) = z_{n+1} | h_t, i_t) \\ &= P(\zeta(h_{t+1}) = z_{n+1} | h_t, j_t) = \frac{1}{2} , \end{aligned}$$

where z_{n+1} stands for either z'_{n+1} or z''_{n+1} , i_t for i' or i'' , and j_t for j' or j'' .

Proof: The choices of i_t and j_t are made independently. Q.E.D.

Thus, in case (ii) a lottery with probabilities $1/2, 1/2$ is performed among z'_{n+1} and z''_{n+1} . Moreover, no player has any control over the outcome--whichever of his two possible moves he chooses, the probabilities are the same ($1/2, 1/2$). Therefore, this is called (following Aumann, Maschler and Stearns [1968]) a jointly controlled lottery.

This completes the definition of the master plan for $t = n!$ (the communication stages). It corresponds to advancing one step in the tree (from z_n to z_{n+1}). We next consider the payoff periods. Let $z_n = \zeta(h_{(n-1)!+1})$ (thus, we are just after z_n was determined at stage $(n-1)!$). Let $\theta_n = \theta_n(z_n)$ in $\Delta^{I \times J}$ be given by Proposition 5.8 (see (5.9)). At stages $(n-1)! + 1$ through $n! - 1$, the players will play θ_n by frequencies; namely, the pair (i, j) will be played $\theta_n(i, j)$ of the time. Since all the denominators are n by (5.12), this can be done in cycles of length n each. For example, assume $\theta_n(i', j') = 1/n$, $\theta_n(i'', j'') = (n-1)/n$ and $\theta_n(i, j) = 0$ otherwise, then player 1 plays i' once (at $t = (n-1)! + 1$), then $n-1$ times i'' (at $t = (n-1)! + 2, \dots, (n-1)! + n$), repeating this n -stage cycle up to (and including) $t = n! - 1$; as for player 2, he chooses j' at $t = (n-1)! + 1$ and j'' at $t = (n-1)! + 2, \dots, (n-1)! + n$, and so on. Clearly, we put $\zeta(h_t) = z_n$ for all $(n-1)! < t \leq n!$, when the two players play as described.

We introduce the following notation: for every i in I , j in J and u, v in N with $u \leq v$, let

$$(5.15) \quad \phi_u^v(i, j) = \frac{1}{v - u + 1} \cdot |\{t \in N : u \leq t \leq v, i_t = i, j_t = j\}|.$$

Thus, $\phi_u^v(i, j)$ is the frequency that the pair of moves (i, j) was used at stages $u, u+1, \dots, v$. Note that it is H_{v+1} -measurable.

Lemma 5.16: Let $t \in N$, $(n-1)! < t < n!$. Then, for all i in I and j in J ,

$$|\phi_{(n-1)!+1}^t(i, j) - \theta_n(i, j)| \leq \frac{n-1}{t - (n-1)!}.$$

Proof: Every n stages, the frequency θ_n is precisely obtained. The inequality follows by ignoring the (at most) $n-1$ stages following the last complete n cycle. Q.E.D.

Finally, we have to define the punishments--what each player does after detecting a deviation from the master plan by the other player. Two results are needed from the theory of zero-sum games (see Propositions 3.5 and 3.6; the more precise statements here are needed to obtain a uniform equilibrium).

Proposition 5.17: Assume the vector y in R^k satisfies (3.3). Then player 2 has a strategy $\bar{\tau} \equiv \bar{\tau}(y)$ such that

$$E_{\sigma', \bar{\tau}}^k(a_T^k) \leq y^k + \frac{2M}{\sqrt{T}},$$

for all strategies σ' of player 1, all k in K and all T in N .

Proof: The precise bound is obtained from the proof of the approachability theorem (cf. Blackwell [1956], or Mertens and Zamir [1980, Ch. I]). Q.E.D.

Proposition 5.18: For every q in Δ^k , player 1 has a strategy $\bar{\sigma} \equiv \bar{\sigma}(q)$ such that

$$E_{\bar{\sigma}, \tau', q}(\beta_T) \leq (\text{vex val}_2 B)(q)$$

for all strategies τ' of player 2 and all T in N .

Proof: The above inequality actually holds with $B^k(i_t, j_t)$ instead of β_T , for all t (e.g., see Mertens and Zamir [1980, Theorem 3.15]). Q.E.D.

The definition of σ and τ can now be completed. Assume first that player 1 deviated from the master plan, either by playing $i_t \neq i^*, i^*$ at some $t = n!$ or by not playing the "right" i_t at some $(n-1)! < t < n!$. Let D be the stage at which this deviation of player 1 occurred. Thus, all moves in h_D are according to the master plan, and i_D is the deviation move (which is observed by player 2 before stage $D+1$). Let $z_n = \zeta(h_D)$ be the corresponding node just before the deviation; the strategy τ prescribes then that after h_{D+1} (i.e., from stage $D+1$ on), player 2 should use $\bar{\tau}(y)$ with $y = f_n(z_n)$ (See Proposition 5.17, and note that (3.3) is satisfied in view of Proposition 3.16(i)).

Next, assume player 2 deviated from the master plan at stage D (and was detected). From stage $D+1$ on, player 1 then uses $\bar{\sigma}(q)$ with $q = p_n(z_n)$ as defined in Proposition 5.18 (again, $z_n = \zeta(h_D)$).

This ends the definition of the pair of strategies σ and τ .

5.4 Payoffs and Probabilities for (σ, τ)

In this subsection we assume that both players use σ and τ , respectively. Thus, only the master plan matters; there are no derivations and no punishments.

We first analyze the payoffs. Let $T \in \mathbb{N}$, $(n-1)! < T \leq n!$, and let h_T in H_T be a history possible under (σ, τ) ; i.e., $P_{\sigma, \tau, p}(h_T) > 0$. We will write θ_n for $\theta_n(\zeta(h_T))$ --the value of θ_n on the atom $\zeta(h_T)$ of Z_n ; similarly for the other random variables defined on Z . Recalling definition (5.15) of the frequencies ϕ , we have

Proposition 5.19: Let $T \in \mathbb{N}$, $(n-1)! < T \leq n!$, $h_T \in H_T$ with $P_{\sigma, \tau, p}(h_T) > 0$. Then, for all i in I and j in J ,

$$|\phi_1^{T-1}(i, j) - [(1 - \frac{(n-1)!}{T-1})\theta_n(i, j) + \frac{(n-1)!}{T-1}\theta_{n-1}(i, j)]| \leq \frac{4}{n}.$$

Proof: If $4/n \geq 1$, there is nothing to prove (both ϕ and the expression [...] lie in the interval $[0, 1]$). Let $n \geq 5$, then

$$\phi_1^{T-1} = \frac{(n-1)!}{T-1} \phi_1^{(n-1)!} + \frac{(T-1) - (n-1)!}{T-1} \phi_1^{(n-1)!+1}.$$

The frequency θ_{n-1} is "played" at stages $(n-2)! < t < (n-1)!$, therefore

$$|\phi_1^{(n-1)!}(i, j) - \theta_{n-1}(i, j)| \leq \frac{(n-2)! + 1}{(n-1)!}$$

(the difference is due to $1 \leq t \leq (n-2)!$ and $t = (n-1)!$; in total, $(n-2)! + 1$ stages; note that θ_{n-1} requires a cycle of length $n-1$, which divides $(n-1)!$). Together with the inequality in Lemma 5.16 for the second term, we obtain an overall difference of at most

$$\frac{(n-2)! + 1 + n - 1}{T - 1} \leq \frac{(n-2)! + n}{(n-1)!} < \frac{4}{n},$$

the last inequality being easily checked (for $n \geq 5$).

Q.E.D.

This result shows that the frequencies obtained are close to those given by the sequence $\{\theta_n\}$. The next two corollaries will be needed in the sequel; we will write $\bar{M} = M \cdot |I| \cdot |J|$.

Corollary 5.20: For all k in K ,

$$a_{T-1}^k \leq f_n^k + \frac{(n-1)!}{T-1} (f_{n-1}^k - f_n^k) + \frac{5\bar{M}}{n}.$$

Proof: By Proposition 5.19,

$$a_{T-1}^k = A^k(\phi_1^{T-1}) \leq \lambda_T A^k(\theta_n) + \lambda'_T A^k(\theta_{n-1}) + \frac{4\bar{M}}{n}$$

with $\lambda'_T = (n-1)!/(T-1)$ and $\lambda_T = 1 - \lambda'_T$.

Recalling (5.10) gives

$$A^k(\theta_n) \leq E(A^k(\theta_\infty) | Z_n) + \frac{\bar{M}}{n}.$$

By (5.6) $A^k(\theta_\infty) \leq r_\infty^k$, hence

$$A^k(\theta_n) \leq r_n^k + \frac{\bar{M}}{n} .$$

Similarly for θ_{n-1} , and we obtain

$$A^k(\phi_1^{T-1}) \leq \lambda_T r_n^k + \lambda_T r_{n-1}^k + \frac{5\bar{M}}{n} ,$$

from which the result follows.

Q.E.D.

Corollary 5.21: For all q in Δ^K ,

$$q \cdot B(\phi_1^{T-1}) \leq \max \{q \cdot B(\theta_n), q \cdot B(\theta_{n-1})\} + \frac{4\bar{M}}{n} .$$

Proof: Immediate from Proposition 5.19.

Q.E.D.

Next we deal with the probability $P_{\sigma, \tau, p}$ and the induced posteriors.

Proposition 5.22: Let $n \in \mathbb{N}$, z_n an atom of Z_n and $T \in \mathbb{N}$, $(n-1)! < T \leq n!$. Then:

$$(5.23) \quad P_{\sigma, \tau, p}(\zeta(h_T) = z_n) = Q(z_n) .$$

$$(5.24) \quad P_{\sigma, \tau, p}(\kappa = k | h_T) = p_n^k(z_n) \text{ for all } h_T \text{ in } H_T \text{ with } \zeta(h_T) = z_n \\ \text{and all } k \text{ in } K.$$

Proof: Induction on n . For $n = 1$, there is only one history h_1 (the "empty" history), thus (5.23) is just $1 = 1$ and (5.24) is $p = p_1$ (recall (5.1)). The induction now proceeds as follows.

At stages $(n-1)! < t < n!$ (payoff stages), neither ζ nor any probabilities change (both players make pure choices). At $t = n!$, the probabilities of $\zeta(h_{t+1}) = z'_{n+1}$ or z''_{n+1} are $1/2$ each by Lemmata 5.13 and 5.14, thus equal to $Q(z'_{n+1}|z_n) = Q(z''_{n+1}|z_n)$ by (5.3) (recall that our G-process is now assumed standard). As for the posterior probabilities (of $\kappa = k$), they change only when there is signalling ($t = n!$ and case (i))--and we use again Lemma 5.13. Q.E.D.

We will now show that (σ, τ) result in the payoffs (a, β) . We need first the following result ($E_{\sigma, \tau}^k$ is the conditional expectation given $\kappa = k$, and $f_n^k = f_n^k(\zeta(h_T))$).

Proposition 5.25: Let $T \in \mathbb{N}$, $(n-1)! < T \leq n!$ and $k \in K$.

Then

$$E_{\sigma, \tau}^k(f_n^k) = a^k.$$

Proof: The probability distribution of $z_n = \zeta(h_T)$ induced by $P_{\sigma, \tau, p}$ is precisely Q (by (5.23)); therefore

$$E_{\sigma, \tau, p}(f_n^k(\zeta(h_T))) = E_{(Z)}(f_n^k),$$

where $E_{(Z)}$ denotes expectation on the space Z (with respect to Q ; note that $E_{\sigma, \tau, p}$ and $E_{\sigma, \tau}^k$ are on Ω). Since $\{f_n^k\}_n$ is a martingale and $f_1 = a$ by (5.1), the above equals a^k .

We claim that the same expectation is obtained when using $P_{\sigma, \tau}^k$ instead of $P_{\sigma, \tau, p}$. Indeed, the induced probability distributions over the tree differ only in case of signalling (in a jointly controlled lottery, it is 1/2 in both cases by Lemma 5.14); however, in that case $f_{m+1} = f_m$, so that the expectation is the same. Q.E.D.

Remark: Actually, the conditional expectations (with respect to P and P^k) are also the same--thus, $\{f_n^k\}$ is a martingale also with respect to the probability distribution induced by P^k . Moreover, any strategy σ' of player 1 that differs from σ only in the probabilities used for signalling has this property (as we shall see in the next subsection).

Proposition 5.26: For all k in K

$$\lim_{T \rightarrow \infty} E_{\sigma, \tau}^k(a_T^k) = a^k,$$

$$\lim_{T \rightarrow \infty} E_{\sigma, \tau, p}(\beta_T) = \beta.$$

Proof: We start with player 2. Let $(n-1)! < T \leq n!$; conditioning on H_T , we obtain ($E = E_{\sigma, \tau, p}$):

$$E(\beta_{T-1} | h_T) = \sum_{k \in K} p_n^k B^k(\phi_1^{T-1}),$$

with $p_n^k = p_n^k(\zeta(h_T))$ as usual. By Proposition 5.19,

$$|B^k(\phi_1^{T-1}) - [\lambda_T B^k(\theta_n) + \lambda'_T B^k(\theta_{n-1})]| \leq \frac{4\bar{M}}{n},$$

where $\lambda'_T = (n-1)/(T-1)$ and $\lambda_T = 1 - \lambda'_T$. As in the proof of Proposition 5.25, the distribution of $z_n = \zeta(h_T)$ induced by $P = P_{\sigma, \tau, p}$ is Q (see (5.23)). Therefore,

$$|E(\beta_{T-1}) - E_{(Z)}(p_n \cdot [\lambda_T B^k(\theta_n) + \lambda'_T B^k(\theta_{n-1})])| \leq \frac{4\bar{M}}{n}.$$

As $n \rightarrow \infty$, $\theta_n \rightarrow \theta_\infty$ and $\theta_{n-1} \rightarrow \theta_\infty$ Q -a.s. by (5.11); also $p_n \rightarrow p_\infty$, hence $(\lambda_T, \lambda'_T \geq 0, \lambda_T + \lambda'_T = 1$ and everything is bounded):

$$\lim_{T \rightarrow \infty} E(\beta_{T-1}) = E_{(Z)}(p_\infty \cdot B(\theta_\infty)).$$

By (5.7), this is $E_{(Z)}(\delta_\infty) = \delta_1 = \beta$ (recall (5.1)).

For player 1, the same argument gives $(\alpha_T$ was defined in (4.16)):

$$\lim_{T \rightarrow \infty} E(\alpha_{T-1}) = E_{(Z)}(p_\infty \cdot A(\theta_\infty)) = E_{(Z)}(p_\infty \cdot f_\infty) = p_1 \cdot f_1 = p \cdot a$$

(see (5.6), Proposition 3.18 and (5.1)). If we condition on κ , we have

$$E(\alpha_{T-1}) = \sum_{\kappa \in K} p^\kappa E^\kappa(a_{T-1}^k),$$

where $E^\kappa = E_{\sigma, \tau}^\kappa = E_{\sigma, \tau, p}(\cdot | \kappa = k)$. As in the proof of Corollary 5.20, we obtain

$$a_{T-1}^k \leq \lambda_T f_n^k + \lambda'_T f_{n-1}^k + \frac{5\bar{M}}{n}.$$

By Proposition 5.25

$$E^k(a_{T-1}^k) \leq \lambda_T a^k + \lambda'_T a^k + \frac{5\bar{M}}{n},$$

hence

$$\limsup_{T \rightarrow \infty} E^k(a_T^k) \leq a^k,$$

which together with

$$\lim_{T \rightarrow \infty} \sum_{k \in K} p^k E^k(a_T^k) = \sum_{k \in K} p^k a^k$$

and $p^k > 0$ for all k completes the proof.

Q.E.D.

Remark: The above proof actually shows that, if both players use (σ, τ) , then the average payoffs converge a.s. to the corresponding $\theta_\infty = \lim_{n \rightarrow \infty} \theta_n$, where $\theta_n = \theta_n(\zeta(h_T))$. Therefore, one can interchange the order of limit as $T \rightarrow \infty$ and expectation (and there is no need for Banach limits!).

At this point we can show intuitively that (σ, τ) is indeed an equilibrium point (at least--in the weak sense (2.7)-(2.8); the complicated inequalities in the next two subsections are in part due to the fact that we want to prove the uniform property (2.11)-(2.12)).

Consider player 1, and fix k in K (the true game). As we noted in the last Remark, the average payoff to player 1 for infinite histories without deviation from the master plan is the corresponding $A^k(\theta_\infty)$, which is $\leq f_\infty^k$. If the game has proceeded up to a point corresponding to the node z_n in the tree, his expected payoff will thus be at most $E^k(f_\infty^k | z_n) = f_n^k$ (see Proposition 5.25 and the Remark following it). This will also be the expected payoff if player 1 decides to deviate now--by Proposition 5.17. This shows that he cannot gain by detectable deviations. How about undetected ones? He can only make those at communication stages (at payoff stages, the moves are pure). If a jointly controlled lottery is performed, he cannot influence the outcome--the two alternatives have probability $1/2$ each no matter what player 1 chooses (since player 2 randomizes truly according to τ). If he is in a signalling case, then $f_{n+1}^k = f_n^k$, thus f_{n+1}^k is constant, and any "signal" he uses gives him the same expectation. Therefore undetectable deviations do not help either, and σ is optimal against τ .

Consider now player 2. Since player 1 uses σ , the posterior probabilities are given by p_n . Therefore the expected average payoff of player 2 at a node z_n --if he does not deviate--is precisely $p_n \cdot E(B(\theta_\infty) | z_n)$, which for n large enough is close to $E(p_\infty \cdot B(\theta_\infty) | z_n) = E(\delta_\infty | z_n) = \delta_n$ (since $p_n \rightarrow p_\infty$). If he makes a detectable deviation, he will get thereafter at most $(\text{vex val}_2 B)(p_n) \leq \delta_n$ --thus he cannot gain by doing so. The only other possible change in strategy is in a

jointly controlled lottery; again, if player 1 uses σ , player 2 cannot influence the resulting probabilities. Thus τ is a best response against σ .

5.5 σ is Optimal Against τ

We will show here that σ is a best response of player 1 against the strategy τ of player 2. Moreover, the uniform condition (2.11) will be proved.

Thus, let $\epsilon > 0$; we have to find $T_0 \equiv T_0(\epsilon)$ such that for all $T > T_0$ and all σ'

$$E_{\sigma', \tau}^k(a_T^k) \leq a^k + \epsilon$$

for all k in K (see (2.13) and Proposition 5.26).

As usual, P, P^k, E and E^k refer to σ, τ, p , whereas P', P'^k, E' and E'^k to σ', τ, p .

If both players use (σ, τ) , there are no deviations from the master plan, and ζ is defined for all possible histories (i.e., those with positive P). However, when we consider alternative strategies, it will be useful to define ζ for all histories (i.e., even those that are not possible under (σ, τ)); $\zeta(h_t)$ will be ζ of the part of h_t up to the first stage a (detectable) deviation occurred. Thus, we define

$$(5.27) \quad D = \sup \{t \in N : P(h_t) > 0\} .$$

D is a random variable on Ω with values in $N \cup \{\infty\}$, and is H_∞ -measurable. For every infinite history, D is the stage of the first detectable deviation, if any; $D = \infty$ otherwise. Note that $P(h_t) > 0$ just means that the sequence of moves used by both players at stages $1, 2, \dots, t-1$ is possible under (σ, τ) ; more precisely, that h_t is possible under (σ, τ) when $\kappa = k$ for some k in K .

Let $DA_t \equiv \min \{D, t\}$, then we define for all t in N and h_t in H_t

$$(5.28) \quad \zeta(h_t) = \zeta(h_{DA_t}) .$$

The right-hand side was given in subsection 5.3; we thus extend the definition of ζ to all histories.

Next, we "translate" the G-process to the space Ω , as follows:

$$(5.29) \quad \hat{g}_t(h_t) \equiv g_m(z_m) ,$$

where $z_m = \zeta(h_t)$ (and thus, by (5.28), we have $(m-1)! < DA_t \leq m!$).

As usual, $\hat{g}_t = (\hat{f}_t, \hat{\delta}_t, \hat{p}_t)$, with $\hat{f}_t = (\hat{f}_t^k)_{k \in K}$ in R^K , $\hat{\delta}_t$ in R and $\hat{p}_t = (\hat{p}_t^k)_{k \in K}$ in Δ^K .

Proposition 5.30: Let $k \in K$. The sequence $\{\hat{f}_t^k\}_{t=1}^\infty$ is a martingale on (Ω, H_∞, P^k) with respect to $\{H_t\}_{t=1}^\infty$. Moreover, for all t in N , h_t in H_t and all i_t in I ,

$$(5.31) \quad E^k(\hat{f}_{t+1}^k | h_t, i_t) = \hat{f}_t^k$$

P^k -a.s. .

This proposition is a crucial assertion in our proof. (5.31) means that the strategies σ and τ have been constructed in such a way that player 1 is indifferent among his various choices of i_t at all histories h_t --and this includes both detectable and undetectable deviations from σ (see also Proposition 5.25 and the subsequent Remark).

Proof: The measurability of \hat{f}_t^k with respect to H_t is immediate by definition. As for the martingale property, namely $E^k(\hat{f}_{t+1}^k | h_t) = \hat{f}_t^k$, it will follow from (5.31) (which is stronger, since it holds for all i_t in I , not only in the average).

We now prove (5.31). It is easy to see by (5.29) that $\hat{f}_{t+1}^k \neq \hat{f}_t^k$ only when $DA_t = m!$ and $DA(t+1) = m! + 1$, and thus $D \geq t+1$ and $t = m!$. Let $z_m = \zeta(h_t)$; since $\hat{f}_{m+1}^k \neq \hat{f}_m^k$, case (ii)--a jointly controlled lottery--occurs at z_m . But i_t must be either i' or i'' (otherwise, player 1 deviated and $D = t$); in both instances, $\zeta(h_{t+1})$ is z'_{m+1} or z''_{m+1} with probability $1/2$ each by Lemma 5.14, and (5.31) reduces to

$$\hat{f}_m^k(z_m) = \frac{1}{2} \hat{f}_{m+1}^k(z'_{m+1}) + \frac{1}{2} \hat{f}_{m+1}^k(z''_{m+1}),$$

which holds by (5.3).

Q.E.D.

For each T in N , we define H_{DAT} to be the finite field generated by all events of the form^{13/}

$$\{h_t \text{ and } DAT \geq t\}$$

for t in N , $t \leq T$ and h_t in H_T . This is the field of events prior to the first detectable deviation. Note that $D + 1$ (but not D) is a stopping time relative to the sequence $\{H_t\}_{t=1}^{\infty}$, and so is $(D + 1) \wedge (T + 1) = (DAT) + 1$. The field of events strictly before $(DAT) + 1$ is precisely H_{DAT} . It is easy to see that $H_{DAT} \subset H_T$, and an atom in H_{DAT} , which we denote by h_{DAT} , is of the form

$$h_{DAT} = \{h_t \text{ and } DAT = t\}$$

for some t in N , $t \leq T$ and h_t in H_t .

From now on, we fix a strategy σ' of player 1, an element k in K , and T in N . To shorten notation, we will write D for DAT and H_D for H_{DAT} .

Consider $E^k(a_{T-1}^k)$; we separate it into three parts: before D , at D , and after D (note that only the first one is always non-empty). Thus,

$$\begin{aligned} a_{T-1}^k &= \frac{D-1}{T-1} a_{D-1}^k + \frac{1}{T-1} A^k(i_D, j_D) \\ &\quad + \frac{1}{T-1} \sum_{t=D+1}^{T-1} A^k(i_t, j_t) . \end{aligned}$$

The middle term is at most $M/(T-1)$, hence

$$\begin{aligned} (5.32) \quad E^k(a_{T-1}^k) &\leq E^k\left(\frac{D-1}{T-1} a_{D-1}^k\right) + \frac{M}{T-1} \\ &\quad + E^k\left(\frac{1}{T-1} \sum_{t=D+1}^{T-1} A^k(i_t, j_t)\right) . \end{aligned}$$

For the first term, we have

Lemma 5.33:

$$E,^k(\frac{D-1}{T-1} a_{D-1}^k) \leq E,^k(\frac{D-1}{T-1} \hat{r}_T^k) + \frac{5\bar{M} + 2M}{n-1}.$$

Proof: We use Corollary 5.20:

$$(5.34) \quad E,^k(\frac{D-1}{T-1} a_{D-1}^k) \leq E,^k(\frac{D-1}{T-1} \hat{r}_m^k) + E,^k(\frac{D-1}{T-1} \frac{(m-1)!}{D-1} (r_{m-1}^k - \hat{r}_m^k)) \\ + 5\bar{M} E,^k(\frac{D-1}{T-1} \frac{1}{m}),$$

where \hat{r}_m^k and r_{m-1}^k are evaluated at the corresponding $\zeta(h_D)$ (note that m here is a random variable, $(m-1)! < D \leq m!$).

By definitions (5.28) and (5.29), the first term is precisely

$$E,^k(\frac{D-1}{T-1} \hat{r}_T^k).$$

The second term is separated into two parts. If $D \leq (n-1)!$, then $m \leq n-1$, hence

$$\frac{(m-1)!}{T-1} \leq \frac{(n-2)!}{(n-1)!} = \frac{1}{n-1},$$

giving a bound of $2M/(n-1)$. Next, we claim that

$$E,^k(\frac{(m-1)!}{T-1} (r_{m-1}^k - \hat{r}_m^k) \cdot \chi_{\{D > (n-1)!\}}) = 0,$$

where $\chi_{\{\dots\}}$ denotes the indicator function of the event $\{\dots\}$.

Let $u \equiv (n-1)!$, then $D > u$ if and only if $P(h_{u+1}) > 0$ (see the definition (5.27) of D). But player 2 does not deviate from τ ,

therefore we have P^k -a.s.: $P(h_{u+1}) > 0$ if and only if $P(h_u, i_u) > 0$.

Conditioning on h_u and i_u gives

$$E^k\left(\frac{u}{T-1} \chi_{\{P(h_u, i_u) > 0\}} \cdot E^k(f_{n-1}^k - f_n^k | h_u, i_u)\right).$$

By (5.29), $f_{n-1}^k = \hat{f}_u^k$ and $f_n^k = \hat{f}_{u+1}^k$; recalling (5.31) shows that the whole expression is zero.

The last term in (5.34) is also separated into two: for $D \leq (n-2)!$,

$$\frac{D-1}{T-1} \cdot \frac{1}{m} \leq \frac{(n-2)!}{(n-1)!} \cdot 1 \leq \frac{1}{n-1};$$

for $D > (n-2)!$, $m \geq n-1$ and

$$\frac{D-1}{T-1} \cdot \frac{1}{m} \leq 1 \cdot \frac{1}{n-1}.$$

This gives a bound of $5\bar{M}/(n-1)$; together with $2M/(n-1)$

from the second term, the proof is completed.

Q.E.D.

For the last term in (5.32), we condition on H_D .

Lemma 5.35:

$$E^k\left(\frac{1}{T-1} \sum_{t=D+1}^{t-1} A^k(i_t, j_t) | H_D\right) \leq \frac{T-D-1}{T-1} \hat{f}_T^k + \frac{2M}{\sqrt{T-1}}.$$

Proof: Given h_D , player 2 uses the punishment strategy $\bar{\tau}(y)$ starting at $t = D + 1$, with $y = f_m(\zeta(h_D)) = \hat{f}_T$ (see (5.29)). The inequality is obtained from Proposition 5.17, applied to

$$\frac{1}{T - D - 1} \sum_{t=D+1}^{T-1} A^k(i_t, j_t) \quad . \quad \text{Q.E.D.}$$

Finally, we have our result.

Proposition 5.36: For every $\varepsilon > 0$ there exists $T_0 \equiv T_0(\varepsilon)$ such that for all $T > T_0$ and all σ'

$$E^k(a_T^k) \leq a^k + \varepsilon$$

for all k in K .

Proof: Combining the inequalities in Lemmata 5.33 and 5.35, we obtain from (5.32)

$$E^k(a_{T-1}^k) \leq E^k\left(\frac{T-2}{T-1} \hat{f}_T^k\right) + \frac{M}{T-1} + \frac{5\bar{M} + 2M}{n-1} + \frac{2M}{\sqrt{T-1}}$$

The first term differs from $E^k(\hat{f}_T^k) = f_1^k = a^k$ (see Proposition 5.30 and (5.1)) by at most another $M/(T-1)$. All additional terms are independent of σ' , and converge to zero as $T \rightarrow \infty$ (hence, $n \rightarrow \infty$ too). Q.E.D.

5.6 τ is Optimal Against σ

Here we prove that τ is a best response of player 2 against player 1 using σ ; as in subsection 5.5, we obtain the uniform property (2.12):

Given $\epsilon > 0$, we show that there is $T_0 \equiv T_0(\epsilon)$ such that for all $T > T_0$ and all τ'

$$E_{\sigma, \tau', p}(\beta_T) \leq \beta + \epsilon$$

(recall Proposition 5.26). We will use the notations E' and P' for σ, τ', p .

In subsection 5.5 the time of the first detectable deviation was defined (see (5.27)); also, the G-process was translated to the space of histories by (5.28) and (5.29). Thus, $\hat{\delta}_t(h_t)$ is the value of the sequence $\{\delta_m(z_m)\}_m$ just before the deviation (if any, up to stage t). We have

Proposition 5.37: The sequence $\{\hat{\delta}_t\}_{t=1}^{\infty}$ is a martingale on (Ω, H_{∞}, P') with respect to $\{H_t\}_{t=1}^{\infty}$.

Proof: Similar to that of Proposition 5.30. Again, we actually prove a stronger assertion (but which will not be needed in the sequel), namely

$$E'(\hat{\delta}_{t+1} | h_t, j_t) = \hat{\delta}_t$$

P' -a.s., for all t in N , h_t in H_t and j_t in J . The only case to check is $t = m!$ and $D \geq t + 1$. The probabilities of z'_{m+1} and z''_{m+1} are $1/2$ each no matter what player 2 chooses at stage t ; this is so by definition of σ in case (i) (signalling), and by Lemma 5.14 in case (ii) (a jointly controlled lottery; since $D \geq t + 1$, $j_t = j'$ or j'' then). Q.E.D.

An important property of σ is that player 2 cannot increase the probability of reaching any node z_n in the tree (he may be able to decrease it by making detectable deviations in previous stages).

Proposition 5.38: Let $t \in N$, $(n-1)! < t \leq n!$, and let z_n be an atom of \mathcal{Z}_n . Then

$$P'(\zeta(h_t) = z_n) \leq Q(z_n) .$$

Proof: Induction on n . For $n = 1$ we clearly have equality. Let $u = (n-1)!$, then definition (5.28) gives

$$\begin{aligned} P'(\zeta(h_t) = z_n) &= P'(\zeta(h_{u+1}) = z_n \text{ and } D \geq u+1) \\ &= P'(\zeta(h_{u+1}) = z_n | \zeta(h_u) = z_{n-1} \text{ and } D \geq u+1) \\ &\quad \cdot P'(D \geq u+1 | \zeta(h_u) = z_{n-1}) \cdot P'(\zeta(h_u) = z_{n-1}) . \end{aligned}$$

The same argument as in the proof of Proposition 5.37 shows that the first factor is $1/2$ (since $D \geq u+1$, in both cases (i) and (ii) the probabilities do not change); the second is at most 1, and the third at most $Q(z_{n-1})$ by induction. This completes the proof, since

$$Q(z_n) = (1/2) \cdot Q(z_{n-1}) \text{ by (5.3).} \quad \text{Q.E.D.}$$

Let τ' be a fixed strategy of player 2, and fix T in N , $(n-1)! < T \leq n!$. As in subsection 5.5, we divide $E'(\beta_{T-1})$ into three parts, as follows (D stands for DAT):

$$(5.39) \quad E'(\beta_{T-1}) \leq E'(\frac{D-1}{T-1} \beta_{D-1}) + \frac{M}{T-1} \\ + E'(\frac{1}{T-1} \sum_{t=D+1}^{T-1} B^k(i_t, j_t)) .$$

In the first term we condition on H_D , to obtain

$$(5.40) \quad E'(\frac{D-1}{T-1} \beta_{T-1}) = E'(\frac{D-1}{T-1} \sum_{k \in K} P'(\kappa = k | H_D) \cdot B^k(\phi_1^{D-1})) .$$

Lemma 5.41: Let h_D be an atom of H_D , and $(m-1)! < D \leq m!$.

Then

$$P'(\kappa = k | h_D) = p_m^k(\zeta(h_D))$$

for all k in K .

Proof: Given h_D , player 1 did not detect any deviations by player 2 up to D . Therefore he played according to the master plan, and the result follows by (5.24). Q.E.D.

Lemma 5.42: For every $\epsilon > 0$ there exists $T_1 \equiv T_1(\epsilon)$ independent of τ' such that

$$E'(\frac{D-1}{T-1} \beta_{T-1}) \leq E'(\frac{D-1}{T-1} \hat{\delta}_T) + \epsilon .$$

Proof: On the space Z , we define for all n in N

$$\eta_n = \max \{p_n \cdot B(\theta_n), p_n \cdot B(\theta_{n-1})\} .$$

Clearly η_n is \mathcal{Z}_n -measurable, and $\eta_n \rightarrow \delta_\infty$ Q-a.s. as $n \rightarrow \infty$ (indeed: $p_n \rightarrow p_\infty$, $\theta_n \rightarrow \theta_\infty$ by (5.11), and $\delta_\infty = p_\infty \cdot B(\theta_\infty)$ by (5.7)). We also have $\delta_n \rightarrow \delta_\infty$ as $n \rightarrow \infty$, therefore

$$\lim_{n \rightarrow \infty} E_{(Z)}(|\eta_n - \delta_n|) = 0$$

(everything is bounded by M). Thus, for every $\varepsilon > 0$ there is $n_1 \equiv n_1(\varepsilon)$ large enough such that

$$E_{(Z)}(|\eta_n - \delta_n|) < \varepsilon$$

for all $n > n_1$.

Using Lemma 5.41 and then Corollary 5.21 (with $q = p_m$), we obtain by (5.40)

$$E'(\frac{D-1}{T-1} \beta_{T-1}) \leq E'(\frac{D-1}{T-1} \eta_m) + 4\bar{M}E'(\frac{D-1}{T-1} \frac{1}{m}),$$

where $(m-1)! < D \leq m!$, and $\eta_m = \eta_m(\zeta(h_D))$. As in the proof of Lemma 5.33, the last term is no more than $4\bar{M}/(n-1)$, which can be made arbitrarily small for large enough T (independent of τ').

Therefore, it remains to bound

$$E'(\frac{D-1}{T-1} |\eta_m - \delta_m|)$$

(recall that $\hat{\delta}_T = \delta_m(\zeta(h_D))$). We separate into three parts: $m \leq n-2$, $m = n-1$, and $m = n$. The first one is bounded by $2M/(n-1)$ (since $D \leq (n-2)!$). Let z_n be an atom of \mathcal{Z}_n , then

$$P'(\zeta(h_D) = z_n) = P'(\zeta(h_T) = z_n) \leq Q(z_n)$$

by (5.28) and Proposition 5.38. This implies that

$$E'(\frac{D-1}{T-1} |\eta_m - \delta_m| \cdot \chi_{\{m=n\}}) \leq E_{(Z)}(|\eta_n - \delta_n|)$$

Similarly, let z_{n-1} be an atom of Z_{n-1} , then

$$P'(\zeta(h_D) = z_{n-1}) = P'(\zeta(h_{n!}) = z_{n-1} \text{ and } D \leq n!) \leq Q(z_{n-1})$$

and

$$E'(\frac{D-1}{T-1} |\eta_m - \delta_m| \cdot \chi_{\{m=n-1\}}) \leq E_{(Z)}(|\eta_{n-1} - \delta_{n-1}|)$$

If $n > n_1(\epsilon)$, both expectations are bounded by ϵ , which completes the proof. Q.E.D.

For the last term in (5.39), we again condition on H_D .

Lemma 5.43:

$$E'(\frac{1}{T-1} \sum_{t=D+1}^{T-1} B^K(i_t, j_t) | H_D) \leq \frac{T-D-1}{T-1} \hat{\delta}_T$$

Proof: By Lemma 5.41, the posteriors at h_D are given by $p_m = p_m(\zeta(h_D))$. From stage $D+1$ on, player 1 uses his punishment strategy; by Proposition 5.18, the expression we consider is thus no more than

$$\frac{T - D - 1}{T - 1} (\text{vex val}_2^B)(p_m) \leq \frac{T - D - 1}{T - 1} \delta_m$$

(we used Proposition 3.16 (ii)). But $\hat{\delta}_T = \delta_m(\zeta(h_D))$, completing the proof. Q.E.D.

Proposition 5.44: For every $\epsilon > 0$ there exists $T_0 \equiv T_0(\epsilon)$ such that for all $T > T_0$ and all τ'

$$E'(\beta_T) \leq \beta + \epsilon.$$

Proof: Similarly to Proposition 5.36, we combine (5.39), Lemmata 5.42 and 5.43, Proposition 5.37, and (5.1). Q.E.D.

We have completed the proof of the second half of our main result.

Proposition 5.45: Let $(a, \beta, p) \in G^*$. Then there exists a uniform equilibrium point (σ, τ) in $\Gamma_\infty(p)$ with payoffs (a, β) .

Proof: Propositions 5.26, 5.36 and 5.44. Q.E.D.

6. Enforceable Joint Plans

Let us consider now equilibria that require finite sequences of communications. For every positive integer m , let

$$G^m = \{g \in G^*: \text{there exists a } G\text{-process } \{g_n\}_{n=1}^\infty \text{ starting at } g \text{ such that } g_n = g_m \text{ for all } n \geq m\}.$$

Thus, G^m corresponds to those G-processes for which the limit g_∞ is reached already at stage m . Clearly, $G^1 = G$ (recall (3.12)). Therefore, the first such set to study is G^2 .

The following is easily obtained: A point $g = (a, \beta, p)$ belongs to G^2 if and only if g can be expressed as a convex combination of points in G , all of which have the same a coordinate or the same p coordinate. Thus, there is a finite set S such that $g = \sum_{s \in S} \rho(s)g(s)$, with $\rho = (\rho(s))_{s \in S}$ in Δ^S , $g(s) = (a(s), \beta(s), p(s))$ in G for all s in S , and either $a(s) = a$ for all s or $p(s) = p$ for all s .

The latter case ($p(s) = p$ for all s) leads to no additional points outside G ; this is due to the fact that, for a fixed p , the set of (a, β) such that (a, β, p) belongs to G is a convex set (indeed, all conditions (3.3), (3.4), (3.9)-(3.11) are invariant under convex combinations--again, when p is constant).

Therefore, the only interesting case is $a(s) = a$ for all s (and $p(s)$ not constant). This generates points in G^2 that do not necessarily belong to G , and that correspond to equilibria with one communication only^{14/} (signalling), followed by payoff accumulation henceforth (using frequencies). Following Aumann, Maschler and Stearns [1968], this is called an enforceable joint plan.^{15/}

An interesting question is: how many different signals are needed? Since the only information player 1 has (that player 2 has not) is the value of κ , it seems reasonable that no more than $|K|$ signals should be required. Namely, the most player 1 can transmit to player 2 is

just κ , which has $|K|$ possible values. However, it turns out that this is not the case, and the correct bound is $|K| + 1$ rather than $|K|$; i.e., no more than $|K| + 1$ signals are needed, and there are examples which do indeed require $|K| + 1$.

For every integer ℓ , let $G^2(\ell)$ be the set of all $g = (a, \beta, p)$ in G^2 such that

$$g = \sum_{s=1}^{\ell} \rho(s) g(s) \quad , \quad \sum_{s=1}^{\ell} \rho(s) = 1 \quad , \quad g(s) = (a, \beta(s), p(s)) \in G$$

and $\rho(s) \geq 0$ for all $s = 1, 2, \dots, \ell$.

Proposition 6.1:

$$G^2 = G^2(|K| + 1) \quad .$$

Proof: For fixed a , the vector (β, p) lies in $R \times \Delta^K$, which is a $|K|$ -dimensional Euclidean space; we now apply Caratheodory's Theorem. Q.E.D.

We will next present an example where $G^2 \neq G^2(|K|)$, showing that $|K| + 1$ is the best bound.

Example 6.2: Let $K = \{1, 2\}$, $I = \{1, 2\}$, $J = \{1, 2, 3, 4, 5, 6, 7\}$.

The two games are (player 1 chooses the row, player 2 the column):

k = 1	$i \backslash j$	1	2	3	4	5	6	7
	1	0,0	0,4	0,-5	-1,-9	-1,-3	-1,3	-1,6
	2	0,0	0,4	0,-5	-1,-9	-1,-3	-1,3	-1,6

k = 2	$i \backslash j$	1	2	3	4	5	6	7
	1	0,0	0,-5	0,4	-1,6	-1,3	-1,-3	-1,-9
	2	0,0	0,-5	0,4	-1,6	-1,3	-1,-3	-1,-9

It is easy to see that $(\text{val}_1 A)(p) = -1$ for all p in Δ^K , and $(\text{vex val}_2 B)(p) = (\text{val}_2 B)(p) = \max \{-9p^1 + 6p^2, -3p^1 + 3p^2, 3p^1 - 3p^2, 6p^1 - 9p^2\}$, where $p = (p^1, p^2)$. Therefore, the intersection of G with the hyperplane $a = (0,0)$ consists of exactly three points:

$$g(1) = ((0,0), 0, (\frac{1}{2}, \frac{1}{2})) ,$$

$$g(2) = ((0,0), 1, (\frac{2}{3}, \frac{1}{3})) ,$$

$$g(3) = ((0,0), 1, (\frac{1}{3}, \frac{2}{3})) ,$$

where we write as usual $g = ((a^1, a^2), \beta, (p^1, p^2))$; these three points correspond to $j = 1$, $j = 2$ and $j = 3$, respectively (i does not matter). Indeed, since $a = (0,0)$, $j = 4, 5, 6$ and 7 are not possible; individual rationality for player 2 (namely (3.4)) then implies that $j = 1$ can be used only at $p = (1/2, 1/2)$, $j = 2$ only at $p = (2/3, 1/3)$, and $j = 3$ only at $p = (1/3, 2/3)$.

Therefore G^2 will contain the convex hull of $g(1), g(2)$ and $g(3)$ --however, no interior point of this triangle can be expressed using only two of its vertices.

It is easily seen in this example that an additional condition may reduce the number of signals to $2 = |K|$. In general, we have

Proposition 6.3: Let $(a, \beta, p) \in G^2$. Then there exists β' in R such that $\beta' \geq \beta$ and $(a, \beta', p) \in G^2(|K|)$.

Proof: By Proposition 6.3,

$$g = (a, \beta, p) = \sum_{s=1}^{\ell} \rho(s) g(s), \quad \sum_{s=1}^{\ell} \rho(s) = 1,$$

$g(s) = (a, \beta(s), p(s)) \in G$ and $\rho(s) \geq 0$ for all $s = 1, 2, \dots, \ell$, where $\ell \leq |K| + 1$. Assume $\ell = |K| + 1$, and consider the ℓ vectors $\{(p(s), 1)\}_{s=1}^{\ell}$ in $\Delta^K \times R$. They must be linearly dependent; let $\{\pi(s)\}_{s=1}^{\ell}$ be not all zero and such that $\sum_{s=1}^{\ell} \pi(s) p(s) = 0$ and $\sum_{s=1}^{\ell} \pi(s) = 0$. Without loss of generality, we assume that $\sum_{s=1}^{\ell} \pi(s) \beta(s) \geq 0$ (otherwise, replace all $\pi(s)$ by $-\pi(s)$). Let $\eta = \min \{-\rho(s)/\pi(s) : \pi(s) > 0\}$, and put $\rho'(s) = \rho(s) + \eta \pi(s)$. Then $\rho'(s) \geq 0$ for all $s = 1, 2, \dots, \ell$ and at least one $\rho'(s)$ is zero; moreover, $\sum_{s=1}^{\ell} \rho'(s) = 1$, $\sum_{s=1}^{\ell} \rho'(s) p(s) = p$ and $\beta' \equiv \sum_{s=1}^{\ell} \rho'(s) \beta(s) \geq \beta$. Q.E.D.

The following is now immediate.

Corollary 6.5: Let (a, β) be the payoffs of a Pareto optimal^{16/} enforceable joint plan equilibrium in $\Gamma_{\infty}(p)$. Then no more than $|K|$ signals are needed; namely, $(a, \beta, p) \in G^2(|K|)$.

What about G^m for larger values of m ? It is easy to see that each G^m is obtained from the previous G^{m-1} by taking convex combinations--with either a fixed (when m is even) or p fixed (when m is odd). New points are usually obtained; Aumann, Maschler and Stearns [1968] provide examples where $G^3 \neq G^2$ and $G^4 \neq G^3$. It is probably not difficult to use the same ideas in order to generate examples where $G^m \neq G^{m-1}$ for arbitrary m .

An open question still remains. Is G^* the union of all the G^m ? If one ignores the game structure, and just considers the notions of bi-convexification (G^m) and bi-martingale (G^*), the answer is negative: G^* may contain points that do not belong to any G^m (and it is not just a matter of closure either-- G^* may be a very different set). For details on these problems, the reader is referred to the forthcoming paper of R. J. Aumann and the author.

Footnotes

- 1/ The standard example is the well known children's way of choosing among two alternatives with equal probability ("two-finger Morra"): they each show, simultaneously, either one or two fingers. If they match (i.e., both show the same number), the first alternative is chosen; if not, the second one. If both choose the number of fingers at random (i.e., with probabilities $1/2, 1/2$), the two alternatives each have probability $1/2$, even when one of the participants uses any other strategy! (This is better than tossing a coin, which may be counterfeited - a fact known to one but not to the other). This idea of jointly controlled randomizations is due to Aumann, Maschler and Stearns [1968].
- 2/ In this case, one may duplicate the single strategy of a player: this enables him to make choices, which do not affect the payoffs, but serve as "signals".
- 3/ The set H_1 , being an empty product, is defined to consist of one element only.
- 4/ A finite field means a field with finitely many elements; such a field is equivalent to a finite partition of the space (the atoms of the field being the elements of the partition).
- 5/ The statement is to be understood as: $[a_\infty \dots]$ if and only if $[a_n \dots]$ for all n ; similarly in (ii).
- 6/ I.e., with δ_∞ or δ_s for β , and p_∞ or p_s for p in (3.4).
- 7/ Henceforth we will always use t for integers in N , and s for half-integers in N_2 .
- 8/ We list here only those we will need in our proofs; the existence of such L is guaranteed by the Hahn-Banach Theorem (see the reference above).
- 9/ I.e., $E(E(Z_{n+1} | F_{n+1}) | F_n) \leq E(Z_n | F_n)$.
- 10/ Z is a set, \mathcal{Z} a σ -field on Z , and Q a probability measure on Z .
- 11/ By "period" we will usually mean the stages from $(n-1)!$ to $n!$ for some n .

- 12/ If there are more than two strategies, the communications may be "shortened" (i.e., less stages). This is not important in our model, since payoffs in finitely many periods do not matter, but will be so if a fixed discount rate is assumed.
- 13/ This is the set of infinite histories which coincide with h_t up to time t , and for which $D \wedge T$ is no less than t .
- 14/ If the G-process is standard (cf. subsection 5.1), this would require one stage in the game; in general, this may take longer (e.g., if player 1 uses only i' and i'' as in subsection 5.3, then at least $\log_2 l$ stages are needed, where l is the number of different values of g_2).
- 15/ They only define "joint plans" - and then find conditions under which these can be "enforced" by equilibria. As Sorin [1981] pointed out, one of their conditions should be slightly strengthened - and then it corresponds to our characterization of G^2 .
- 16/ I.e., such that there is no other enforceable joint plan equilibrium in $\Gamma(p)$ with payoffs (a', β') satisfying $(a', \beta') \geq (a, \beta)$ and $(a', \beta') \neq (a, \beta)$.

References

- Aumann, R.J. [1964], "Mixed and Behaviour Strategies in Infinite Extensive Games", in: Advances in Game Theory (Annals of Mathematics Studies 52), M. Dresher et al (eds.), Princeton University Press, Princeton, 627-650.
- Aumann, R.J. [1976], "Agreeing to Disagree", Annals of Statistics 4, 1236-1239.
- Aumann, R.J. [1981], "Repetition as a Paradigm for Cooperation in Games of Incomplete Information", mimeo.
- Aumann, R.J. and M. Maschler [1966], "Game Theoretic Aspects of Gradual Disarmament", in: Mathematica [1966], V.1 - V.55.
- Aumann, R.J., M. Maschler and R.E. Stearns [1968], "Repeated Games of Incomplete Information: An Approach to the Non-Zero-Sum Case", in: Mathematica [1968], 117-216.
- Blackwell, D. [1956], "An Analogue of the Minimax Theorem for Vector Payoffs", Pacific Journal of Mathematics 6, 1-8.
- Dunford, N. and J.T. Schwartz [1958], Linear Operators, Vol. I, Interscience, New York.
- Harsanyi, J. [1967-1968], "Games with Incomplete Information Played by Bayesian Players", Parts I, II, III, Management Science 14, 159-182, 320-334, 486-502.
- Hart, S. [1979], "Lecture Notes on Special Topics in Game Theory", IMSSS-Economics, Stanford University, mimeo.
- Hildenbrand, W. [1974], Core and Equilibria in a Large Economy, Princeton University Press, Princeton.
- Kuhn, H.W. [1953], "Extensive Games and the Problem of Information", in: Contributions to the Theory of Games II (Annals of Mathematics Studies 28), H.W. Kuhn and A.W. Tucker (eds.), Princeton University Press, Princeton, 193-216.
- Mathematica [1966-1968], Reports to the U.S. Arms Control and Disarmament Agency, prepared by Mathematica, Inc., Princeton. ST-80 [1966], ST-116 [1967], ST-140 [1968].

- Mertens, J.-F. and S. Zamir [1980], "Repeated Games", mimeo.
- Meyer, P.A. [1973], "Limites Mediales d'apres Mokobodzki", in: *Seminaire de Probabilites VIII*, Universite de Strasbourg (1971/1972), Springer-Verlag, 198-204.
- Sorin, S. [1980], "An Introduction to Two-Person Zero-Sum Repeated Games with Incomplete Information", IMSSS - Economics TR-312, Stanford University.
- Sorin, S. [1981], "Some Results on the Existence of Nash Equilibria for Non-Zero Sum Games with Incomplete Information", DP-8107, CORE.

REPORTS IN THIS SERIES

160. "The Structure and Stability of Competitive Dynamical Systems," by David Cass and Karl Shell.
161. "Monopolistic Competition and the Capital Market," by J. E. Stiglitz.
162. "The Corporation Tax," by J. F. Stiglitz.
163. "Measuring Returns to Scale in the Aggregate and the Scale Effect of Public Goods," by David A. Starrett.
164. "Monopoly, Quality, and Regulation," by Michael Spence.
165. "A Note on the Belief Constraint in a Model of Borrowing," by Duncan K. Foley and Martin F. Hellwig.
166. "Incentives, Risk, and Information: Notes Towards a Theory of Hierarchy," by Joseph E. Stiglitz.
167. "Asymptotic Expansions of the Distributions of Estimates in Simultaneous Equations for Alternative Parameter Sequences," by T. W. Anderson.
168. "Estimation of Linear Functional Relationships: Approximate Distributions and Connections with Simultaneous Equations in Econometrics," by T. W. Anderson.
169. "Monopoly and the Rate of Extraction of Exhaustible Resources," by Joseph E. Stiglitz.
170. "Equilibrium in Competitive Insurance Markets: An Essay on the Economics of Imperfect Information," by Michael Rothschild and Joseph Stiglitz.
171. "Strong Consistency of Least Squares Estimates in Normal Linear Regression," by T. W. Anderson and John B. Taylor.
172. "Incentive Schemes under Differential Information Structures: An Application to Trade Policy," by Partha Dasgupta and Joseph Stiglitz.
173. "The Incidence and Efficiency Effects of Taxes on Income from Capital," by John B. Shoven.
174. "Distribution of a Maximum Likelihood Estimate of a Slope Coefficient: The LIML Estimate for Known Covariance Matrix," by T. W. Anderson and Takamitsu Sawa.
175. "A Comment on the Test of Overidentifying Restrictions," by Joseph B. Kadane and T. W. Anderson.
176. "An Asymptotic Expansion of the Distribution of the Maximum Likelihood Estimate of the Slope Coefficient in a Linear Functional Relationship," by T. W. Anderson.
177. "Some Experimental Results on the Statistical Properties of Least Squares Estimates in Control Problems," by T. W. Anderson and John B. Taylor.
178. "A Note on 'Fulfilled Expectations' Equilibria," by David M. Kreps.
179. "Uncertainty and the Rate of Extraction under Alternative Institutional Arrangements," by Partha Dasgupta and Joseph E. Stiglitz.
180. "Budget Displacement Effects of Inflationary Finance," by Jerry Green and E. Sheshinski.
181. "Towards a Marxist Theory of Money," by Duncan K. Foley.
182. "The Existence of Futures Markets, Noisy Rational Expectations and Informational Externalities," by Sanford Grossman.
183. "On the Efficiency of Competitive Stock Markets where Traders have Diverse Information," by Sanford Grossman.
184. "A Bidding Model of Perfect Competition," by Robert Wilson.
185. "A Bayesian Approach to the Production of Information and Learning by Doing," by Sanford J. Grossman, Richard E. Kihlstrom and Leon J. Mirman.
186. "Disequilibrium Allocations and Recontracting," by Jean-Michel Grandmont, Guy Laroque and Yves Younes.
187. "Agreeing to Disagree," by Robert J. Aumann.
188. "The Maximum Likelihood and the Nonlinear Three Stage Least Squares Estimator in the General Nonlinear Simultaneous Equation Model," by Takeshi Anemiyama.
189. "The Modified Second Round Estimator in the General Qualitative Response Model," by Takeshi Anemiyama.
190. "Some Theorems in the Linear Probability Model," by Takeshi Anemiyama.
191. "The Bilinear Complementarity Problem and Competitive Equilibria of Linear Economic Models," by Robert Wilson.
192. "Noncooperative Equilibrium Concepts for Oligopoly Theory," by L. A. Gerard-Varet.
193. "Inflation and Costs of Price Adjustment," by Eytan Sheshinski and Yoram Weiss.
194. "Power and Taxes in a Multi-Community Economy," by R. J. Aumann and M. Kurz.
195. "Distortion of Preferences, Income Distribution and the Case for a Linear Income Tax," by Mordecai Kurz.
196. "Search Strategies for Nonrenewable Resource Deposits," by Richard J. Gilbert.
197. "Demand for 'Irrational' Factors, Inflation and Adjustment Costs," by Eytan Sheshinski and Yoram Weiss.
198. "Bargains and Ripoffs: A Model of Monopolistically Competitive Price Dispersions," by Steve Salop and Joseph Stiglitz.
199. "The Design of Tax Structure: Direct Versus Indirect Taxation by A. B. Atkinson and J. E. Stiglitz.
200. "Market Allocations of Location Choice in a Model with Free Mobility," by David Starrett.
201. "Efficiency in the Optimum Supply of Public Goods," by Lawrence J. Lau, Eytan Sheshinski and Joseph E. Stiglitz.
202. "Risk Sharing, Sharecropping and Uncertain Labor Markets," by David M. G. Newberry.
203. "On Non-Walrasian Equilibria," by Frank Hahn.
204. "A Note on Elasticity of Substitution Functions," by Lawrence J. Lau.
205. "Quantity Constraints as Substitutes for Inoperative Markets: The Case of the Credit Markets," by Mordecai Kurz.
206. "Incremental Consumer's Surplus and Helonic Price Adjustment," by Robert D. Willig.
207. "Optimal Depletion of an Uncertain Stock," by Richard Gilbert.
208. "Some Minimum Chi-Square Estimators and Comparisons of Normal and Logistical Models in Qualitative Response Analysis," by Kunio Morimune.
209. "A Characterization of the Optimality of Equilibrium in Incomplete Markets," by Sanford J. Grossman.
210. "Inflation and Taxes in a Growing Economy with Debt and Equity Finance," by M. Feldstein, J. Green and E. Sheshinski.
211. "The Specification and Estimation of a Multivariate Logit Model," by Takeshi Anemiyama.
212. "Prices and Queues as Screening Devices in Competitive Markets," by Joseph E. Stiglitz.
213. "Conditions for Strong Consistency of Least Squares Estimates in Linear Models," by T. W. Anderson and John B. Taylor.
214. "Utilitarianism and Horizontal Equity: The Case for Random Taxation," by Joseph E. Stiglitz.
215. "Simple Formulas for Optimal Income Taxation and the Measurement of Inequality," by David M. Kreps and Eytan Sheshinski.
216. "Temporal Resolution of Uncertainty and Dynamic Choice Behavior," by David M. Kreps and Eytan Sheshinski.
217. "The Estimation of Nonlinear Labor Supply Functions with Taxes from a Truncated Sample," by Michael Hurd.
218. "The Welfare Implications of the Unemployment Rate," by Michael Hurd.
219. "Keynesian Economics and General Equilibrium Theory: Reflections on Some Current Debates," by Frank Hahn.
220. "The Core of an Exchange Economy with Differential Information," by Robert Wilson.
221. "A Competitive Model of Exchange," by Robert Wilson.
222. "Intermediate Preferences and the Majority Rule," by Jean-Michel Grandmont.
223. "The Fixed Price Equilibria: Some Results in Local Comparative Statics," by Guy Laroque.
224. "On Stockholder Unanimity in Making Production and Financial Decisions," by Sanford J. Grossman and Joseph E. Stiglitz.
225. "Selection of Regressors," by Takeshi Anemiyama.
226. "A Note on a Heteroscedastic Model," by Takeshi Anemiyama.
227. "Welfare Measurement for Local Public Finance," by David Starrett.
228. "Unemployment Equilibrium with Rational Expectations," by W. P. Heller and R. M. Starr.
229. "A Theory of Competitive Equilibrium in Stock Market Economies," by Sanford J. Grossman and Oliver D. Hart.
230. "An Application of Stein's Methods to the Problem of Single Period Control of Regression Models," by Asad Zaman.
231. "Second Best Welfare Economics in the Mixed Economy," by David Starrett.
232. "The Logic of the Fix-Price Method," by Jean-Michel Grandmont.
233. "Tables of the Distribution of the Maximum Likelihood Estimate of the Slope Coefficient and Approximations," by T. W. Anderson and Takamitsu Sawa.
234. "Further Results on the Informational Efficiency of Competitive Stock Markets," by Sanford Grossman.
235. "The Estimation of a Simultaneous-Equation Tobit Model," by Takeshi Anemiyama.
236. "The Estimation of a Simultaneous-Equation Generalized Probit Model," by Takeshi Anemiyama.
237. "The Consistency of the Maximum Likelihood Estimator in a Disequilibrium Model," by T. Anemiyama and G. Sen.
238. "Numerical Evaluation of the Exact and Approximate Distribution Functions of the Two-Stage Least Squares Estimate," by T. W. Anderson and Takamitsu Sawa.
239. "Risk Measurement of Public Projects," by Robert Wilson.
240. "On the Capitalization Hypothesis in Closed Communities," by David Starrett.
241. "A Note on the Uniqueness of the Representation of Commodity-Augmenting Technical Change," by Lawrence J. Lau.
242. "The Property Rights Doctrine and Demand Revelation under Incomplete Information," by Kenneth J. Arrow.
243. "Optimal Capital Gains Taxation Under Limited Information," by Jerry R. Green and Eytan Sheshinski.
244. "Straightforward Individual Incentive Compatibility in Large Economies," by Peter J. Hammond.
245. "On the Rate of Convergence of the Core," by Robert J. Aumann.
246. "Unsatisfactory Equilibria," by Frank Hahn.
247. "Existence Conditions for Aggregate Demand Functions: The Case of a Single Index," by Lawrence J. Lau.
248. "Existence Conditions for Aggregate Demand Functions: The Case of Multiple Indices," by Lawrence J. Lau.
249. "A Note on Exact Index Numbers," by Lawrence J. Lau.
250. "Linear Regression Using Both Temporally Aggregated and Temporally Disaggregated Data," by Cheng Hsiao.
251. "The Existence of Economic Equilibria: Continuity and Mixed Strategies," by Partha Dasgupta and Eric Maskin.
252. "A Complete Class Theorem for the Control Problem and Further Results on Admissibility and Inadmissibility," by Asad Zaman.
253. "Measure-Based Values of Market Games," by Sergiu Hart.
254. "Altruism as an Outcome of Social Interaction," by Mordecai Kurz.
255. "A Representation Theorem for 'Preference for Flexibility'," by David M. Kreps.
256. "The Existence of Efficient and Incentive Compatible Equilibria with Public Goods," by Theodore Groves and John O. Ledyard.
257. "Efficient Collective Choice with Compensation," by Theodore Groves.
258. "On the Impossibility of Informationally Efficient Markets," by Sanford J. Grossman and Joseph E. Stiglitz.
- 259.

REPORTS IN THIS SERIES

- 260 "Values for Games Without Subgame Perfect Equilibria: Some Difficulties With Current Concepts," by Alvin E. Roth
- 261 "Martingales and the Valuation of Redundant Assets," by J. Michael Harrison and David M. Kreps
- 262 "Autoregressive Modelling and Money Income Causality Detection," by Cheng Hsiao
- 263 "Measurement Error in a Dynamic Simultaneous Equations Model Without Stationary Disturbances," by Cheng Hsiao
- 264 "The Measurement of Deadweight Loss Revisited," by W. E. Diewert
- 265 "The Elasticity of Derived Net Supply and a Generalized Le Chatelier Principle," by W. E. Diewert
- 266 "Income Distribution and Distribution of Preferences: The E-Commodity Case," by Mordchai Kurz
- 267 "The a-Order Mean Squared Errors of the Maximum Likelihood and the Minimum Logit Chi-Square Estimator," by Takeshi Anemiyu
- 268 "Temporal Variations in the Distribution of the Maximum Likelihood and the Minimum Logit Chi-Square Estimator," by Takeshi Anemiyu
- 269 "Take-Over Bids and the Theory of the Corporation," by David M. Kreps and Evan L. Porteus
- 270 "The Numerical Values of Some Key Parameters in Econometric Models," by T. W. Anderson, Kimo Morimune, and Takamitsu Sawa
- 271 "Less Representations of Information Structures and their Comparisons," by Jerry Green and Nancy Stokey
- 272 "Asymptotic Expansions of the Distributions of Estimators in a Linear Functional Relationship when the Sample Size is Large," by Naoto Kunimoto
- 273 "Public Goods and Power," by R. J. Aumann, M. Kurz, and A. Neyman
- 274 "An Asymptotic Approach to the Efficiency of Non-Cooperative Equilibrium in Economies with a Continuum of Traders," by A. Mas-Colell
- 275 "Tables of the Exact Distribution Function of the Limited Information Maximum Likelihood Estimator when the Covariance Matrix is Known," by T. W. Anderson and Takamitsu Sawa
- 276 "Autoregressive Modelling of Canadian Money and Income Data," by Cheng Hsiao
- 277 "We Can't Disagree Forever," by John D. Gensink and Heraklis Polemarchakis
- 278 "Constrained Excess Demand Functions," by Heraklis Polemarchakis
- 279 "On the Bayesian Solution of Nash Equilibrium," by Akira Futukida
- 280 "Disequilibrium Economics in Simultaneous Equations Systems," by C. Gourinchon, J. J. Laffont, and A. Monfort
- 281 "Duality Approaches to Microeconomic Theory," by W. E. Diewert
- 282 "A Time Series Analysis of the Impact of Canadian Wage and Price Controls," by Cheng Hsiao and Oluwalayo Fakayea
- 283 "A Strategic Theory of Inflation," by Mordchai Kurz
- 284 "A Characterization of Vector Measure Games in \mathbb{R}^n ," by Yair Lurman
- 285 "On the Method of Exclusion and the Provision of Local Public Goods," by David A. Starrett
- 286 "An Optimization Problem Arising in Economics: Approximate Solutions, Linearities, and a Law of Large Numbers," by Sergiu Hart
- 287 "Asymptotic Expansions of the Distributions of the Estimates of Coefficients in a Simultaneous Equation System," by Yasuhiro Fukuda, Kimo Morimune, Naoto Kunimoto, and Masahiko Taniguchi
- 288 "Optimal Voluntary Income Distribution," by R. J. Arrow
- 289 "Asymptotic Values of Mixed Games," by Abraham Neyman
- 290 "Time Series Modelling and Causal Ordering of Canadian Money, Income and Interest Rate," by Cheng Hsiao
- 291 "An Analysis of Power in Exchange Economies," by Martin J. Osborne
- 292 "Estimation of the Regression of a Normal Mean," by Asad Zaman
- 293 "Improving the Maximum Likelihood Estimate in Linear Functional Relationships for Alternative Parameter Sequences," by Kimo Morimune and Naoto Kunimoto
- 294 "Calculation of Bivariate Normal Integrals by the Use of Incomplete Negative-Order Moments," by Kei Takeuchi and Akimichi Akemichi
- 295 "On Partitioning of a Sample with Binary Type Questions in Lieu of Collecting Observations," by Kenneth J. Arrow, Leon Pechuchinsky and Milton Sobel
- 296 "The Two-Stage Least Squares Deviations Estimation," by Takeshi Anemiyu
- 297 "Three Essays on Capital Markets," by David M. Kreps
- 298 "Infinite Horizon Programs," by Michael J. P. Magill
- 299 "Electoral Outcomes and Social Log Likelihood Maxima," by Peter Coughlin and Shmuel Nitzan
- 300 "Notes on Social Choice and Voting," by Peter Coughlin
- 301 "Overlapping Expectations and Hart's Conditions for Equilibrium in a Securities Model," by Peter J. Hammond
- 302 "Directional and Local Electorate Competitions with Probabilistic Voting," by Peter Coughlin and Shmuel Nitzan
- 303 "Asymptotic Expansions of the Distributions of the Test Statistics for Overidentifying Restrictions in a System of Simultaneous Equations," by Kimo Morimune, and Takamitsu Sawa
- 304 "Incomplete Markets and the Observability of Risk Preference Properties," by H. H. Polemarchakis and L. Selden
- 305 "Multi-Period Securities and the Efficient Allocation of Risk: A Comment on the Black-Scholes Option Pricing Model," by David M. Kreps
- 306 "Asymptotic Expansions of the Distributions of k-Class Estimators when the Disturbances are Small," by Naoto Kunimoto, Kimo Morimune, and Yoshiko Fukuda
- 307 "Arbitrage and Equilibrium in Economies with Infinitely Many Commodities," by David M. Kreps
- 308 "Unemployment Equilibrium in an Economy with Linked Prices," by Mordchai Kurz
- 309 "Pareto Optimal Nash Equilibria are Competitive in a Repeated Economy," by Mordchai Kurz and Sergiu Hart
- 310 "Identification," by Cheng Hsiao
- 311 "An Introduction to Two Person Zero-Sum Repeated Games with Incomplete Information," by Sylvain Sorin
- 312 "Estimation of Dynamic Models With Error Components," by T. W. Anderson and Cheng Hsiao
- 313 "On Robust Estimation in Certainty Equivalence Control," by Anders R. Westlund and Hans Stenlund
- 314 "The Social Costs of Monopoly and Regulation: A Game Theoretic Analysis," by William P. Rogerson
- 315 "On Industry Equilibrium Under Uncertainty," by J. Drèze and E. Sheshinski
- 316 "Cost Benefit Analysis and Project Evaluation From the Viewpoint of Productive Efficiency," by W. E. Diewert
- 317 "On the Chain-Store Paradox and Predation: Reputation for Toughness," by David M. Kreps and Robert Wilson
- 318 "On the Number of Commodities Required to Represent a Market Game," by Sergiu Hart
- 319 "Evaluation of the Distribution Function of the Limited Information Maximum Likelihood Estimator," by T. W. Anderson, Naoto Kunimoto, and Takamitsu Sawa
- 320 "A Comparison of the Logit Model and Normal Discriminant Analysis When the Independent Variables Are Binary," by Takeshi Anemiyu and James L. Powell
- 321 "Efficiency of Resource Allocation by Uninformed Demand," by Theodore Groves and Sergiu Hart
- 322 "A Comparison of the Box-Cox Maximum Likelihood Estimator and the Nonlinear Two-Stage Least Squares Estimator," by Takeshi Anemiyu and James L. Powell
- 323 "Comparison of the Densities of the TSLS and LIML Estimators for Simultaneous Equations," by T. W. Anderson, Naoto Kunimoto, and Takamitsu Sawa
- 324 "Admissibility of the Bayes Procedure Corresponding to the Uniform Prior Distribution for the Control Problem in Four Dimensions but Not in Five," by Charles Stein and Asad Zaman
- 325 "Some Recent Developments on the Distributions of Single-Equation Estimators," by T. W. Anderson
- 326 "On Inflation," by Frank Hahn
- 327 "Two Papers on Majority Rule: 'Continuity Properties of Majority Rule with Intermediate Preferences,' by Peter Coughlin and Shmuel Nitzan; and, 'Electoral Outcomes with Probabilistic Voting and Nash Social Welfare Maxima,' by Peter Coughlin and Shmuel Nitzan
- 328 "On the Endogenous Formation of Coalitions," by Sergiu Hart and Mordchai Kurz
- 329 "Controllability, Pecuniary Externalities and Optimal Taxation," by David Starrett
- 330 "Nonlinear Regression Models," by Takeshi Anemiyu
- 331 "Paradoxical Results From Inada's Conditions for Majority Rule," by Herve Raynaud
- 332 "On Welfare Economics with Incomplete Information and the Social Value of Public Information," by Peter J. Hammond
- 333 "Equilibrium Policy Proposals With Abstentions," by Peter J. Coughlin
- 334 "Infinite Excessive and Invariant Measures," by Michael I. Takser
- 335 "The Life-Cycle Hypothesis and the Effects of Social Security and Private Pensions on Family Savings," by Mordchai Kurz
- 336 "Optimal Retirement Age," by Mordchai Kurz
- 337 "Bayesian Incentive Compatible Beliefs," by Claude d'Aspremont and Louis-Andre Gerard-Varet
- 338 "Qualitative Response Models: A Survey," by Takeshi Anemiyu
- 339 "The Social Costs of Monopoly and Regulation: A Game Theoretic Analysis," by William P. Rogerson

AD-A123 587

NON-ZERO-SUM TWO-PERSON REPEATED GAMES WITH INCOMPLETE
INFORMATION(U) STANFORD UNIV CA INST FOR MATHEMATICAL
STUDIES IN THE SOCIAL SCIENCES S HART FEB 82 TR-367
N00014-79-C-0685

2/2

UNCLASSIFIED

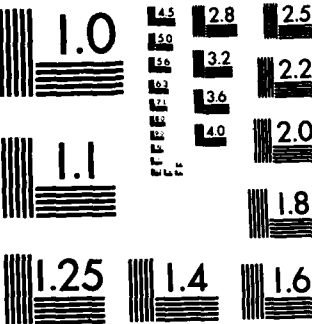
F/G 5/11

NL

END

FORMED

DATE



MICROCOPY RESOLUTION TEST CHART
NATIONAL BUREAU OF STANDARDS-1963-A

REPORTS IN THIS SERIES

260. "Values for Games Without Sidepayments: Some Difficulties With Current Concepts," by Abin E. Roth.
261. "Hartung and the Valuation of Redundant Assets," by J. Michael Harrison and David M. Kreps.
262. "Autoregressive Modelling and Money Income Causality Detection," by Cheng Hsiao.
263. "Measurement Error in a Dynamic Simultaneous Equations Model without Stationary Disturbances," by Cheng Hsiao.
264. "The Measurement of Deadweight Loss Revisited," by W. E. Diewert.
265. "The Elasticity of Derived Net Supply and a Generalized Le Chatelier Principle," by W. E. Diewert.
266. "Income Distribution and Distribution of Preferences: The R Commodity Case," by Mordecai Kurz.
267. "The π^2 Order Mean Squared Errors of the Maximum Likelihood and the Minimum Logit Chi-Square Estimator," by Takeshi Amemiya.
268. "Temporal Von Neumann-Morgenstern and Induced Preferences," by David M. Kreps and Evan I. Porteus.
269. "Take-Over Bids and the Theory of the Corporation," by Stanford Grossman and Oliver D. Hart.
270. "The Numerical Values of Some Key Parameters in Econometric Models," by T. W. Anderson, Kimio Morimune and Takamitsu Sawa.
271. "Two Representations of Information Structures and their Comparisons," by Jerry Green and Nancy Stokey.
272. "Asymptotic Expansions of the Distributions of Estimators in a Linear Functional Relationship when the Sample Size is Large," by Naoto Kunitomo.
273. "Public Goods and Power," by R. J. Aumann, M. Kurz and A. Neyman.
274. "An Asymptotic Approach to the Efficiency of Non-Cooperative Equilibrium in Economies with a Continuum of Traders," by A. Mas-Colell.
275. "Tables of the Exact Distribution Function of the Limited Information Maximum Likelihood Estimator when the Covariance Matrix is Known," by T. W. Anderson and Takamitsu Sawa.
276. "Autoregressive Modeling of Canadian Money and Income Data," by Cheng Hsiao.
277. "We Can't Disagree Forever," by John D. Gauthier and Heraklis Polemarchakis.
278. "Constrained Excess Demand Functions," by Heraklis M. Polemarchakis.
279. "On the Bayesian Selection of Nash Equilibrium," by Akira Tomioka.
280. "Disequilibrium Economics in Simultaneous Equations Systems," by C. Gourieroux, J. J. Laffont and A. Monfort.
281. "Duality Approaches to Microeconomic Theory," by W. E. Diewert.
282. "A Time Series Analysis of the Impact of Canadian Wage and Price Controls," by Cheng Hsiao and Oluwasayo Fajana.
283. "A Strategic Theory of Inflation," by Mordecai Kurz.
284. "A Characterization of Vector Measure Games in gNA ," by Yari Tsutsui.
285. "On the Method of Taxation and the Provision of Local Public Goods," by David A. Starrett.
286. "An Optimization Problem Arising in Economics: Approximate Solutions, Linearity, and a Law of Large Numbers," by Sergiu Hart.
287. "Asymptotic Expansions of the Distributions of the Estimators of Coefficients in a Simultaneous Equation System," by Yasunori Kunitomo, Naoto Kunitomo and Hiroshi Taniguchi.
288. "Optimal & Voluntary Income Distribution," by K. J. Arrow.
289. "Asymptotic Values of Mixed Games," by Abraham Neyman.
290. "Time Series Modelling and Causal Ordering of Canadian Money, Income and Interest Rate," by Cheng Hsiao.
291. "An Analysis of Power in Exchange Economies," by Martin J. Osborne.
292. "Formation of the Recognition of a Normal Mean," by Asad Zaman.
293. "Improving the Maximum Likelihood Estimate in Linear Functional Relationships for Alternative Parameter Sequences," by Kimio Morimune and Naoto Kunitomo.
294. "Calculation of Bivariate Normal Integrals by the Use of Incomplete Negative-Order Moments," by Kei Takeshi and Akimichi Takemura.
295. "On Partitioning of a Sample with Binary-Type Questions in Lieu of Collecting Observations," by Kenneth J. Arrow, Leon Pechinskiy and Milton Sobel.
296. "The Two-Sample Least Absolute Deviations Estimator," by Takeshi Amemiya.
297. "Three Essays on Capital Markets," by David M. Kreps.
298. "Infinite Horizon Programs," by Michael J. P. Magill.
299. "Electoral Outcomes and Social Log-Likelihood Maxima," by Peter Coughlin and Shmuel Nitzan.
300. "Notes on Social Choice and Voting," by Peter Coughlin.
301. "Overlapping Expectations and Hart's Conditions for Equilibrium in a Securities Model," by Peter J. Hammond.
302. "Directional and Local Electorate Competitions with Probabilistic Voting," by Peter Coughlin and Shmuel Nitzan.
303. "Asymptotic Expansions of the Distributions of the Test Statistics for Overidentifying Restrictions in a System of Simultaneous Equations," by Kimimori Morimune and Takada.
304. "Incomplete Markets and the Observability of Risk Preference Properties," by H. H. Polemarchakis and L. Selden.
305. "Multiperiod Securities and the Efficient Allocation of Risk: A Comment on the Black-Scholes Option Pricing Model," by David M. Kreps.
306. "Asymptotic Expansions of the Distributions of k-Class Estimators when the Disturbances are Small," by Naoto Kunitomo, Kimio Morimune, and Yoshiko Takada.
307. "Arbitrage and Equilibrium in Economies with Infinitely Many Commodities," by David M. Kreps.
308. "Unemployment Equilibrium in an Economy with Linked Prices," by Mordecai Kurz.
309. "Pareto Optimal Nash Equilibria are Competitive in a Repeated Economy," by Mordecai Kurz and Sergiu Hart.
310. "Identification," by Cheng Hsiao.
311. "An Introduction to Two-Person Zero Sum Repeated Games with Incomplete Information," by Sylvain Sorin.
312. "Estimation of Dynamic Models With Error Components," by T. W.
313. "On Robust Estimation in Certainty Equivalence Control," by Anders H. Nestlund and Hans Stenlund.
314. "On Industry Equilibrium Under Uncertainty," by J. Dreze and Z. Szehinski.
315. "Cost Benefit Analysis and Project Evaluation From the Viewpoint of Productive Efficiency," by W. E. Diewert.
316. "On the Chain-Store Paradox and Predation: Reputation for Toughness," by David M. Kreps and Robert Wilson.
317. "On the Number of Commodities Required to Represent a Market Game," by Sergiu Hart.
318. "Evaluation of the Distribution Function of the Limited Information Maximum Likelihood Estimator," by T. W. Anderson, Naoto Kunitomo, and Takamitsu Sawa.
319. "A Comparison of the Logit Model and Normal Discriminant Analysis When the Independent Variables Are Binary," by Takeshi Amemiya and James L. Powell.
320. "Efficiency of Resource Allocation by Uninformed Demand," by Theodore Groves and Sergiu Hart.
321. "A Comparison of the Box-Cox Maximum Likelihood Estimator and the Nonlinear Two Stage Least Squares Estimator," by Takeshi Amemiya and James L. Powell.
322. "Comparison of the Densities of the TSLS and LIML Estimators for Simultaneous Equations," by T. W. Anderson, Naoto Kunitomo, and Takamitsu Sawa.
323. "Admissibility of the Bayes Procedure Corresponding to the Uniform Prior Distribution for the Control Problem in Four Dimensions but Not in Five," by Charles Stein and Asad Zaman.
324. "Some Recent Developments on the Distributions of Single-Equation Estimators," by T. W. Anderson.
325. "On Inflation," by Frank Hahn.
326. "Two Papers on Majority Rule: 'Continuity Properties of Majority Rule with Intermediate Preferences,' by Peter Coughlin and Kuan-Pin Lin; and 'Electoral Outcomes with Probabilistic Voting and Nash Social Welfare Maxima,' by Peter Coughlin and Shmuel Nitzan.
327. "On the Endogenous Formation of Coalitions," by Sergiu Hart and Mordecai Kurz.
328. "Controllability, Pecuniary Externalities and Optimal Taxation," by David Starrett.
329. "Nonlinear Regression Models," by Takeshi Amemiya.
330. "Paradoxical Results From Inada's Conditions for Majority Rule," by Hervé Raynaud.
331. "On Welfare Economics with Incomplete Information and the Social Value of Public Information," by Peter J. Hammond.
332. "Equilibrium Policy Proposals With Abstentions," by Peter J. Coughlin.
333. "Infinite Excessive and Invariant Measures," by Michael I. Takser.
334. "The Life-Cycle Hypothesis and the Effects of Social Security and Private Pensions on Family Savings," by Mordecai Kurz.
335. "Optimal Retirement Age," by Mordecai Kurz.
336. "Bayesian Incentive Compatible Beliefs," by Claude d'Aspremont and Louis-André Gérard-Varet.
337. "Qualitative Response Models: A Survey," by Takeshi Amemiya.
338. "The Social Costs of Monopoly and Regulation: A Game Theoretic Analysis," by William P. Rogerson.

Reports in this Series

- 340. "Sequential Equilibria," by David M. Kreps and Robert Wilson.
- 341. "Enhancing of Semigroups," by Michael I. Taksar.
- 342. "Formulation and Estimation of Dynamic Models Using Panel Data," by T.W. Anderson and Cheng Hsiao.
- 343. "Ex-Post Optimality as a Dynamically Consistent Objective for Collective Choice Under Uncertainty," by Peter Hammond.
- 344. "Three Lectures In Monetary Theory," by Frank H. Hahn.
- 345. "Socially Optimal Investment Rules in the Presence of Incomplete Markets and Other Second Best Distortions," by Frank Milne and David A. Starrett.
- 346. "Approximate Purification of Mixed Strategies," by Robert Aumann, Yitzhak Katznelson, Roy Radner, Robert W. Rosenthal, and Benjamin Weiss.
- 347. "Conditions for Transitivity of Majority Rule with Algorithmic Interpretations," by Herve Raynaud.
- 348. "How Restrictive Actually are the Value Restriction Conditions," by Herve Raynaud.
- 349. "Cournot Duopoly in the Style of Fulfilled Expectations Equilibrium," by William Novshek and Hugo Sonnenschein.
- 350. "Law of Large Numbers for Random Sets and Allocation Processes," by Zvi Artstein and Sergiu Hart.
- 351. "Risk Perception in Psychology and Economics," by Kenneth J. Arrow.
- 352. "Shrunken Predictors for Autoregressive Models," by Taku Yamamoto
- 353. "Predation, Reputation, and Entry Deterrence," by Paul Milgrom and John Roberts.
- 354. "Social and Private Production Objectives in the Sequence Economy" by David Starrett
- 355. "Recursive Rational Choice" by Alain Lewis
- 356. "Least Absolute Deviations Estimation for Censored and Truncated Regression Models" by James Powell
- 357. "Relatively Recursive Rational Choice" by Alain Lewis
- 358. "A Theory of Auctions and Competitive Bidding" by Paul Milgrom and Robert Weber.
- 359. "Facing an Uncertain Future" by W. M. Gorman
- 360. "The Individual Freedom Allowed by the Value Restriction Condition" by H. Raynaud

Reports in this Series

- 361. "Incomplete Resource Allocation Mechanisms" by P. J. Hammond
- 362. "The Comparative Dynamics of Efficient Programs of Capital Accumulation and Resource Depletion" by W. E. Diewert and T. R. Lewis
- 363. "Optimum Pricing Policy Under Stochastic Inflation" by E. Sheshinski and Yoram Weiss
- 364. "Capital Accumulation and the Characteristics of Private Intergenerational Transfers" by M. Kurz
- 365. "Asymptotic Efficiency and Higher Order Efficiency of the Limited Information Maximum Likelihood Estimator in Large Econometric Models" by Naoto Kunitomo
- 366. "Semi-Values of Political Economic Games," by Abraham Neyman
- 377. "Non-Zero-Sum Two-Person Repeated Games with Incomplete Information," by Sergiu Hart